arXiv:1602.08807v1 [stat.ME] 29 Feb 2016

Exploratory data analysis for extreme values using non-parametric kernel methods

B. Béranger*†‡, T. Duong†§, S. E. Perkins-Kirkpatrick¶, and S. A. Sisson‡

1 March 2016

Abstract

In many settings it is critical to accurately model the extreme tail behaviour of a random process. Non-parametric density estimation methods are commonly implemented as exploratory data analysis techniques for this purpose as they possess excellent visualisation properties, and can naturally avoid the model specification biases implied by using parametric estimators. In particular, kernel-based estimators place minimal assumptions on the data, and provide improved visualisation over scatterplots and histograms. However kernel density estimators are known to perform poorly when estimating extreme tail behaviour, which is important when interest is in process behaviour above some large threshold, and they can over-emphasise bumps in the density for heavy tailed data. In this article we develop a transformation kernel density estimator, and demonstrate that its mean integrated squared error (MISE) efficiency is equivalent to that of standard, non-tail focused kernel density estimators. Estimator performance is illustrated in numerical studies, and in an expanded analysis of the ability of well known global climate models to reproduce observed temperature extremes in Sydney, Australia.

Keywords: Exploratory data analysis; Extreme value theory; Global climate models; Histograms; Multivariate kernel density estimation; Model selection.

1 Introduction

The extreme values (very large or very small values) of a dataset are frequently of interest as they are closely related to uncommon events with important consequences. For climate data, these extreme events include heat waves (prolonged extreme high temperatures), cold snaps (extreme low temperatures), floods (extreme high levels of waterways or tides or waves), storms (extreme high wind speeds or amounts of precipitation) and droughts (prolonged extreme low amounts of rainfall) (e.g. Kotz and Nadarajah, 2000, Coles, 2001).

Suppose that $\boldsymbol{X} = (X_1, \ldots, X_d)^\top$ is a d-dimensional random vector with cumulative distribution function (c.d.f.) $F_{\boldsymbol{X}}$ and probability density function (p.d.f.) $f_{\boldsymbol{X}}$. A common representation of the extremes arising from this distribution focuses on examining those

*Corresponding author. Email: b.beranger@unsw.edu.au

†Theoretical and Applied Statistics Laboratory (LSTA) University Pierre and Marie Curie - Paris 6, F-75005, Paris, France

‡School of Mathematics and Statistics, University of New South Wales, Sydney, Australia

§Current address: Computer Science Laboratory (LIPN) University Paris-Nord - Paris 13, F-93430, Villetaneuse, France

¶Climate Change Research Centre, University of New South Wales, Sydney, Australia

values which exceed some high threshold $\boldsymbol{u} = (u_1, \ldots, u_d)^\top$, which determines the support on which values of $\boldsymbol{X}$ are considered to be extreme. We can denote these extreme values by $\boldsymbol{X}^{[\boldsymbol{u}]} \equiv \boldsymbol{X}|\boldsymbol{X} > \boldsymbol{u}$, under which each marginal inequality must hold i.e. $X_j > u_j$ for $j = 1, \ldots, d$. Estimating the tail behaviour of $\boldsymbol{X}$ is one of the goals of extreme value theory.

In the simplest case of univariate $\boldsymbol{X}$ ($d = 1$) extremes, common approaches to distributional tail estimation rely on parametric models, typically based on generalised extreme value (GEV) or generalised Pareto distributions (GPD), for which numerous estimation procedures are available. These methods include maximum likelihood (Prescott and Walden, 1980, Hosking, 1985, Smith, 1985, Macleod, 1989), probability weighted moments (Hosking et al., 1985), maximum product spacing (Cheng and Amin, 1983), least squares estimation (Maritz and Munro, 1967), estimation based on order statistics and records (Pickands, 1975, Hill, 1975), the method of moments (Christopeit, 1994) and Bayesian estimation (Lye et al., 1993). As with all parametric estimators, these approaches suffer from the possibility of misspecification, particularly when the asymptotic GEV and GPD models can't be assumed to hold. This potential can be avoided by non-parametric estimation which does not make assumptions on particular parametric forms. See e.g. Markovich (2007, Chapter 3) for a summary of non-parametric estimation of univariate heavy-tail densities.

For multivariate extremes ($d > 1$), non-parametric estimation of indicators of extremal dependence is an intensively studied field, and includes estimation of the Pickands or extremal dependence function (Pickands, 1975, Hall and Tajvidi, 2000, Marcon et al., 2014), the tail dependence function (Huang, 1992, Drees and Huang, 1998, Einmahl et al., 2008, 2012), and the spectral measure (Einmahl et al., 2001, Einmahl and Segers, 2009, de Carvalho et al., 2013). The motivation for non-parametric estimators is stronger for multivariate extremes than in the univariate case as there is no general parametric form to describe the range of extremal behaviour of max-stable processes. This means that a choice of any particular parametric family has a high chance of resulting in model misspecification.

In this article we propose a new class of multivariate kernel-based nonparametric density estimators to model the extreme tail behaviour of $\boldsymbol{X}$. Kernel density estimators are among the most widely used non-parametric estimators (e.g. Wand and Jones, 1995, Simonoff, 1996) and they possess excellent visualisation properties which can naturally form part of an exploratory data analysis. However, standard kernel estimators can produce spurious bumps in the estimated tails of $f_{\boldsymbol{X}}$ if it has heavy tails, as is particularly the case when analysing extremes. The standard estimators are also strongly influenced by edge effects due to truncated support, which is again of particular interest in modelling the extremes when focus is on $\boldsymbol{X}^{[\boldsymbol{u}]}$, which is truncated at the threshold $\boldsymbol{u}$. Here we focus on modifications to standard kernel density estimation which attenuate these spurious bumps, and accommodates a truncated support when estimating the tail density $f_{\boldsymbol{X}^{[\boldsymbol{u}]}}$.

Our approach focuses on direct density estimation of $f_{\boldsymbol{X}^{[\boldsymbol{u}]}}$ rather than of $f_{\boldsymbol{X}}$, which will tend to reduce the appearance of spurious bumps in the density estimate as a kernel smoothing parameter more appropriate for the tail region can be chosen. The problem of truncated support is handled by constructing a transformation of the domain of $\boldsymbol{X}^{[\boldsymbol{u}]}$ to $\mathbb{R}^d$ and then implementing standard kernel density estimation techniques in the transformed space (see e.g. Charpentier and Flachaire, 2014, Silverman, 1986). In this article we establish the mean integrated squared error (MISE) optimality of this implied approach, and demonstrate that the performance of the transformation kernel density estimator is asymptotically as good as the performance of standard kernel density estimators on non-

bounded spaces. In the context of extremes, this permits the construction and theoretical justification of more efficient kernel based nonparametric density estimators for the tails of observed processes. It additionally allows these estimators to be used within existing goodness-of-fit measures (e.g. Perkins et al., 2013) in place of more poorly performing histogram estimates of tail behaviour.

The layout of this article is as follows. Our primary contribution is presented in Section 2, which develops the transformation kernel density estimator, establishes its asymptotic (MISE) efficiency (with proofs deferred to the Appendix) and examines optimal bandwidth estimation. We also assess histogram based tail density estimation, and develop the role of tail density estimators in goodness-of-fit (model selection) procedures. In Section 3 we verify our results on finite samples for simulated data in both univariate and multivariate settings, and in Section 4 we expand the work of Perkins et al. (2013) by performing an analysis of 22 global climate models (GCMs) and assess how well they are able to reproduce observed temperature extremes. Section 5 concludes with a discussion.

2 Tail densities for extreme values

2.1 Transformation tail density estimation

Let $X_1, \ldots, X_m$ be a random sample drawn from a common univariate distribution F_X with density f_X. Let n of these observations exceed the threshold u, to form the tail sample $X_1^{[u]}, \ldots, X_n^{[u]}$ where $X_i^{[u]} \in (u, \infty), i = 1, \ldots, n$. Standard kernel estimators are most often defined for data with unbounded support. Using these kernel estimators without modification for data on (u, ∞) with substantial amounts of data close to u can lead to increased bias at the boundary of the support. One major class of modified estimators are based on the modifying the kernel function itself to avoid assigning probability mass outside the data support (e.g. the linear boundary kernel (Gasser and Müller, 1979) or the boundary beta kernel (Chen, 1999). However, we instead focus on an alternative modification based on transformation kernel estimators, where a known monotonic transformation $t(\cdot)$ maps the data support to the real line where standard kernel estimators are well-established, before back-transforming to the original bounded support. See e.g. Devroye and Györfi (1985), Silverman (1986), Wand and Jones (1995).

Let $Y = t(X^{[u]})$ be a transformed random variable, with distribution F_Y and density f_Y. The relationship between the transformed random variable Y and the original $X^{[u]}$ at a non-random point x is given by

$$f_{X^{[u]}}(x) = |t'(x)| f_Y(t(x))$$

where t' is the first derivative of t. Rather than directly approaching the difficult task of estimating $f_{X^{[u]}}$, we approach it indirectly via the simpler task of estimating f_Y. Consider the transformed sample $Y_1, \ldots, Y_n$ where $Y_i = t(X_i), i = 1, \ldots, n$, and $y = t(x)$. Since $X_1^{[u]}, \ldots X_n^{[u]}$ are supported on (u, ∞), one suitable transformation would be $t(x) = \log(x - u)$. For data-based choices of the transformation function see Ruppert and Wand (1992) and Ruppert and Cline (1994). As $Y_1, \ldots, Y_n$ are supported on the real line, f_Y can then be estimated by the standard kernel density estimator

$$\hat{f}_Y(y; h) = n^{-1} \sum_{i=1}^{n} K_h(y - Y_i)$$

where $K_h(y) = h^{-1} K(y/h)$ is a scaled kernel, $h > 0$ is the bandwidth or smoothing parameter and K is a symmetric kernel density function. The tail density estimator for

$f_{X^{[u]}}$ can then be defined by replacing the true density by its kernel estimator

$$\hat{f}_{X^{[u]}}(x;h) = |t'(t^{-1}(y))|\hat{f}_Y(y;h).$$

A generalisation of this transformation kernel estimator to multivariate data is established through a d-dimensional random vector $\boldsymbol{X} = (X_1, \ldots, X_d)^\top$ with distribution function $F_{\boldsymbol{X}}$ and density function $f_{\boldsymbol{X}}$. The random variable of values greater than a vector threshold $\boldsymbol{u} = (u_1, \ldots, u_d)^\top$ is denoted as $\boldsymbol{X}^{[u]} \equiv \boldsymbol{X}|\boldsymbol{X} > \boldsymbol{u}$ under which each marginal inequality must hold, i.e. $X_j > u_j$ for $j = 1, \ldots, d$. The support of $\boldsymbol{X}^{[\boldsymbol{u}]}$ is the Cartesian product $(\boldsymbol{u}, \infty) = (u_1, \infty) \times \cdots \times (u_d, \infty)$. For $\boldsymbol{x} \in (\boldsymbol{u}, \infty)$, the corresponding tail density is $f_{\boldsymbol{X}^{[u]}}(\boldsymbol{x}) = f_{\boldsymbol{X}}(\boldsymbol{x})/\bar{F}_{\boldsymbol{X}}(\boldsymbol{u})$ and tail distribution is $F_{\boldsymbol{X}^{[u]}}(\boldsymbol{x}) = F_{\boldsymbol{X}}(\boldsymbol{x})/\bar{F}_{\boldsymbol{X}}(\boldsymbol{u})$, where $\bar{F}_{\boldsymbol{X}}(\boldsymbol{u}) = \int_{(\boldsymbol{u},\infty)} f_{\boldsymbol{X}}(\boldsymbol{w})d\boldsymbol{w}$ is the survival function of $\boldsymbol{X}$ evaluated at $\boldsymbol{u}$.

If $\boldsymbol{X}_1, \ldots, \boldsymbol{X}_m$ form a random sample drawn from the common d-variate distribution $F_{\boldsymbol{X}}$, then suppose n of these observations exceed the threshold $\boldsymbol{u}$, to form the tail sample $\boldsymbol{X}_1^{[\boldsymbol{u}]}, \ldots, \boldsymbol{X}_n^{[\boldsymbol{u}]}$. Consider the transformed random variable $\boldsymbol{Y} = \boldsymbol{t}(\boldsymbol{X}^{[\boldsymbol{u}]})$ where $\boldsymbol{t} : (\boldsymbol{u}, \infty) \to \mathbb{R}^d$ is defined by $\boldsymbol{t}(\boldsymbol{x}) = (t_1(x_1), \ldots, t_d(x_d))^\top$ where the t_j are monotonic functions on (u_j, ∞) e.g. $t_j(x_j) = \log(x_j - u_j), j = 1, \ldots, d$. The density of $\boldsymbol{X}^{[\boldsymbol{u}]}$ is then related to the density of $\boldsymbol{Y}$ by

$$f_{\boldsymbol{X}^{[u]}}(\boldsymbol{x}) = f_{\boldsymbol{Y}}(\boldsymbol{t}(\boldsymbol{x}))|\mathbf{J}_{\boldsymbol{t}}(\boldsymbol{x})|$$

where $|\mathbf{J}_{\boldsymbol{t}}|$ is the Jacobian of $\boldsymbol{t}$. Denoting the transformed data sample as $\boldsymbol{Y}_1, \ldots, \boldsymbol{Y}_n$, with $\boldsymbol{Y}_i = \boldsymbol{t}(\boldsymbol{X}_i), i = 1, \ldots, n$, the kernel estimator of $f_{\boldsymbol{Y}}$ at a non-random point $\boldsymbol{y} = (y_1, \ldots, y_d)^\top = \boldsymbol{t}(\boldsymbol{x})$ is then given by

$$\hat{f}_{\boldsymbol{Y}}(\boldsymbol{y}; \mathbf{H}) = n^{-1} \sum_{i=1}^{n} K_{\mathbf{H}}(\boldsymbol{y} - \boldsymbol{Y}_i)$$

where K is a symmetric d-variate density function, the bandwidth matrix $\mathbf{H}$ is a $d \times d$ positive definite symmetric matrix of smoothing parameters, and the scaled kernel $K_{\mathbf{H}}(\boldsymbol{y}) = |\mathbf{H}^{-1/2}|K(\mathbf{H}^{-1/2}\boldsymbol{y})$. The tail density can then be defined by replacing the true density function by its kernel estimator

$$\hat{f}_{\boldsymbol{X}^{[u]}}(\boldsymbol{x}; \mathbf{H}) = |\mathbf{J}_{\boldsymbol{t}}(\boldsymbol{t}^{-1}(\boldsymbol{y}))|\hat{f}_{\boldsymbol{Y}}(\boldsymbol{y}; \mathbf{H}) \tag{1}$$

where $\boldsymbol{t}^{-1}(\boldsymbol{y}) = (t_1^{-1}(y_1), \ldots, t_d^{-1}(y_d))^\top$ is the element-wise inverse of $\boldsymbol{t}(\boldsymbol{y})$.

2.2 Tail density estimator performance

Under the regularity conditions (A1)–(A3) in the Appendix, the mean integrated squared error (MISE) of the density estimator $\hat{f}_{\boldsymbol{Y}}$ is

$$\operatorname{MISE} \hat{f}_{\boldsymbol{Y}}(\cdot; \mathbf{H}) = \mathbb{E} \int_{\mathbb{R}^d} [\hat{f}_{\boldsymbol{Y}}(\boldsymbol{y}; \mathbf{H}) - f(\boldsymbol{y})]^2 \mathrm{d}\boldsymbol{y} = O(\operatorname{tr} \mathbf{H}^2 + n^{-1}|\mathbf{H}|^{-1/2})$$

using established procedures (see Lemma 1 in the Appendix). That the MISE of $\hat{f}_{\boldsymbol{X}^{[u]}}$ is the same order follows from immediately from Equation (1). Using these MISE expressions, we can then define optimal bandwidth choice as the minimiser of the MISE

$$\mathbf{H}^* = \operatorname*{argmin}_{\mathbf{H} \in \mathcal{F}} \operatorname{MISE} \hat{f}_{\boldsymbol{Y}}(\cdot; \mathbf{H}) = O(n^{-2/(d+4)}) \tag{2}$$

where $\mathcal{F}$ is the space of $d \times d$ symmetric positive definite matrices.

The following theorem asserts that in optimally estimating $\hat{f}_{\boldsymbol{X}^{[\boldsymbol{u}]}}$ we can carry out all the important steps of bandwidth selection and estimation for the transformed data $\boldsymbol{Y}$ with unbounded support, allowing us to take advantage of existing results and algorithms. Moreover, the transformation $\boldsymbol{t}$ does not affect the asymptotic behaviour of the density estimator $\hat{f}_{\boldsymbol{X}^{[\boldsymbol{u}]}}$, and so we are able to recover the same standard optimal convergence rates ($O(n^{-4/(d+4)})$) of the unbounded data estimator.

Theorem 1. *Suppose that the regularity conditions (A1)–(A4) hold (see Appendix). The minimal MISE rate of the tail density estimator $\hat{f}_{\boldsymbol{X}^{[\boldsymbol{u}]}}$ is asymptotically the same order as the minimal MISE rate of the transformed kernel density estimator $\hat{f}_{\boldsymbol{Y}}$, as $n \to \infty$. I.e.*

$$\inf_{\mathbf{H}\in\mathcal{F}} \operatorname{MISE} \hat{f}_{\boldsymbol{X}^{[\boldsymbol{u}]}}(\cdot;\mathbf{H}) - \left\{ \inf_{\mathbf{H}\in\mathcal{F}} \operatorname{MISE} \hat{f}_{\boldsymbol{Y}}(\cdot;\mathbf{H}) \right\} = O(n^{-4/(d+4)}).$$

Proof in the Appendix.

However, if we contrast this performance to that obtained by computing a standard kernel estimator of type $\hat{f}_{\boldsymbol{Y}}$ to the untransformed bounded tail sample $\boldsymbol{X}_1^{[\boldsymbol{u}]}, \ldots, \boldsymbol{X}_n^{[\boldsymbol{u}]}$, the MISE is $O(n^{-1}|\mathbf{H}|^{-1/2}+\operatorname{tr}\mathbf{H})$ which leads to a minimal MISE rate of $O(n^{-2/(d+2)})$ which is asymptotically slower than the minimal MISE rates of $O(n^{-4/(d+4)})$ in Theorem 1 (see e.g. Marshall and Hazelton, 2010). This implies that it is important to take the boundedness of the tail sample into account.

2.3 Optimal bandwidth computation

The optimal bandwidth selector defined in Equation (2) is mathematically intractable as it depends on unknown quantities. Accordingly a vast body of research in the density estimation literature has focused on providing data-based bandwidth selectors which estimate or approximate the optimal bandwidth. There are three main classes: (i) normal scale (or rule of thumb), (ii) plug-in and (iii) cross validation.

The class of normal scale selectors is an extension to the multivariate case of the quick and simple bandwidth selectors where the unknown density f is replaced by a normal density, leading to

$$\hat{\mathbf{H}}_{\mathrm{NS}} = \left[\frac{4}{(d+2)n}\right]^{2/(d+4)} \mathbf{S} n^{-2/(d+4)}$$

where $\mathbf{S}$ is the sample covariance matrix of $\boldsymbol{Y}_1, \ldots, \boldsymbol{Y}_n$ (see e.g. Wand and Jones, 1995, p. 111).

The class of plug-in selectors consists of a generalisation of the work of Sheather and Jones (1991) by Wand and Jones (1994) and Duong and Hazelton (2003). Plug-in selectors use as a starting point the AMISE formula (Asymptotic MISE) where the only unknown quantity is the $\boldsymbol{\psi}_{\boldsymbol{Y},4} = \int_{\mathbb{R}^d} \mathsf{D}^{\otimes 4} f_{\boldsymbol{Y}}(\boldsymbol{y}) f_{\boldsymbol{Y}}(\boldsymbol{y}) \mathrm{d}\boldsymbol{y}$ functional, as defined in Chacón and Duong (2010). The fourth order differential $\mathsf{D}^{\otimes 4}$ is expressed as a vector of length d^4, resulting from a four-fold Kronecker product of the first order differential D, see Holmquist (1996). Replacing this by an estimator $\hat{\boldsymbol{\psi}}_{\boldsymbol{Y},4}$ yields the plug-in criterion

$$\mathrm{PI}(\mathbf{H}) = \tfrac{1}{4} m_2^2(K) (\operatorname{vec}^\top \mathbf{H} \otimes \operatorname{vec}^\top \mathbf{H}) \hat{\boldsymbol{\psi}}_{\boldsymbol{Y},4}(\mathbf{G}) + n^{-1} R(K) |\mathbf{H}|^{-1/2}$$

where $m_2(K)$ is defined in condition (A2) in the Appendix, vec is the operator that stacks the element of a matrix column-wise into a vector, $\hat{\boldsymbol{\psi}}_{\boldsymbol{Y},4}(\mathbf{G}) = n^{-2} \sum_{i,j=1}^n \mathsf{D}^{\otimes 4} L_{\mathbf{G}}(\boldsymbol{Y}_i - \boldsymbol{Y}_j)$, $L_{\mathbf{G}}$ is an initial pilot kernel with pilot bandwidth matrix $\mathbf{G}$ and $R(K) = \int_{\mathbb{R}^d} K(\boldsymbol{x})^2 \mathrm{d}\boldsymbol{x}$. The plug-in selector $\hat{\mathbf{H}}_{\mathrm{PI}}$ is the minimiser over $\mathcal{F}$ of $\mathrm{PI}(\mathbf{H})$.

For the class of cross validation selectors we focus on unbiased (or least squares) cross validation and smoothed cross validation. Unbiased cross validation (UCV) was introduced by Bowman et al. (1984) and Rudemo (1982). The unbiased cross validation selector, $\hat{\mathbf{H}}_{\mathrm{UCV}}$, is defined as the minimiser over $\mathcal{F}$ of

$$\mathrm{UCV}(\mathbf{H}) = \int_{\mathbb{R}^d} \hat{f}_{\boldsymbol{Y}}(\boldsymbol{y}; \mathbf{H})^2 \mathrm{d}\boldsymbol{x} - 2n^{-1} \sum_{i=1}^{n} \hat{f}_{\boldsymbol{Y},-i}(\boldsymbol{Y}_i; \mathbf{H}),$$

where $\hat{f}_{\boldsymbol{Y},-i}(\boldsymbol{Y}_i; \mathbf{H}) = [n(n-1)]^{-1} \sum_{j=1}^{n} K_{\mathbf{H}}(\boldsymbol{Y}_i - \boldsymbol{Y}_j)$. The smoothed cross validation (SCV) selector $\hat{\mathbf{H}}_{\mathrm{SCV}}$, is defined as the minimiser over $\mathcal{F}$ of

$$\begin{aligned}\mathrm{SCV}(\mathbf{H}) = n^{-2} \sum_{i=1}^{n} \sum_{j=1}^{n} (K_{\mathbf{H}} * K_{\mathbf{H}} * L_{\mathbf{G}} * L_{\mathbf{G}} - 2K_{\mathbf{H}} * L_{\mathbf{G}} * L_{\mathbf{G}} + L_{\mathbf{G}} * L_{\mathbf{G}})(\boldsymbol{Y}_i - \boldsymbol{Y}_j) \\ + n^{-1} R(K) |\mathbf{H}|^{-1/2},\end{aligned}$$

where $*$ is the convolution operator (Hall et al., 1992). If there are no replications in the data, then SCV with $\mathbf{G} = 0$ is identical to UCV as the pilot kernel L_0 can then be thought of as the Dirac delta function.

The UCV selector can be directly computed as it contains no unknown quantities, however specification of the bandwidth $\boldsymbol{G}$ of the pilot kernel is required for the plug-in and SCV selectors. Computational data-based algorithms which address this are Wand and Jones (1995), Duong and Hazelton (2003) for plug-in selectors and Hall et al. (1992), Duong and Hazelton (2005) for SCV selectors.

2.4 Tail density estimation via histograms

Histograms, especially for univariate data, are widely used as alternatives to kernel estimators for visualising data samples, even when focusing on distributional tails (see e.g. Perkins et al., 2007, 2013). Their advantages include computational and mathematical simplicity, and that they do not suffer from the boundary bias problems of standard kernel estimators. In the context of tail density estimation, we divide the data range of the sample $\boldsymbol{X}_1^{[\boldsymbol{u}]}, \ldots, \boldsymbol{X}_n^{[\boldsymbol{u}]}$ into a regular partition of hypercubes A_i of size $b_1 \times \cdots \times b_d$, and define the binwidth as $\boldsymbol{b} = (b_1, \ldots, b_d)^\top \in \mathbb{R}^d$. The histogram estimator of $\tilde{f}_{\boldsymbol{X}^{[\boldsymbol{u}]}}$ at a point $\boldsymbol{x}$ in a bin A_i is

$$\tilde{f}_{\boldsymbol{X}^{[\boldsymbol{u}]}}(\boldsymbol{x}; \boldsymbol{b}) = \frac{\gamma_i}{n b_1 \cdots b_d}$$

where γ_i represents the number of observations in the hypercube A_i. If conditions similar to (A1) and (A3) in the Appendix are fulfilled then, by Scott (2015, Theorem 3.5), the MISE of the histogram estimator is

$$\mathrm{MISE}\, \tilde{f}_{\boldsymbol{X}^{[\boldsymbol{u}]}}(\cdot; \boldsymbol{b}) = O((n b_1 \cdots b_d)^{-1} + \boldsymbol{b}^\top \boldsymbol{b})$$

with minimal MISE

$$\inf_{\boldsymbol{b} > 0} \mathrm{MISE}\, \tilde{f}_{\boldsymbol{X}^{[\boldsymbol{u}]}}(\cdot; \boldsymbol{b}) = O(n^{-2/(d+2)}).$$

where $\boldsymbol{b} > 0$ means that each $b_i > 0$, $i = 1, \ldots, d$. This is is asymptotically slower than the $O(n^{-4/(d+4)})$ minimal MISE rate for the transformation kernel estimator $\tilde{f}_{\boldsymbol{X}^{[\boldsymbol{u}]}}$ from Theorem 1. Hence from the MISE perspective, the transformation kernel estimator is preferable to a histogram for tail density estimation, especially as the dimension d increases.

In analogy with the data-based optimal bandwidth selectors, the normal scale optimal binwidth (Scott, 2015, Theorem 3.5) is

$$\hat{b}_j = 2 \times 3^{1/(d+2)} \pi^{d/(d+4)} s_j n^{-1/(d+2)} \tag{3}$$

where $s_j, j = 1, \ldots, d$ are the marginal sample standard deviations of $\boldsymbol{X}_1^{[\boldsymbol{u}]}, \ldots, \boldsymbol{X}_n^{[\boldsymbol{u}]}$. There is no equivalent variety of binwidth selectors which generalise Equation (3) compared to bandwidth selectors (Section 2.3) due the slower asymptotic performance of histograms as compared to kernel estimators.

2.5 Model assessment via tail density estimation

Tail density estimation can provide one way to assess the fidelity of the observed dataset to one or more candidate models. For example, in climate science different climate models commonly produce competing predictions of environmental variables. The performance of these models is often validated by comparing the model predicted output, with that of the observed data (Flato et al., 2013). These comparisons may be based on the full body of the predicted variables, or focus primarily on the extremes (e.g. Perkins et al., 2007, 2013). Similarly, in the context of extreme value theory, the analyst is regularly required to determine which of multiple competing parametric families, such as max-stable distributions, provides the best fit to an observed extremal dataset (e.g. Coles and Tawn, 1994)

Suppose that we have a suite of parametric models indexed by $\mathcal{M} = \{1, \ldots, M\}$, and we wish to determine which of them most appropriately describe the tails of the underlying distribution of the observed dataset, $f_{\boldsymbol{X}}$. Perkins et al. (2013) ulitised the histogram estimator $\tilde{f}_{\boldsymbol{X}^{[u]}}$ of the observed data sample as a surrogate for the unknown target $f_{\boldsymbol{X}^{[u]}}$, and so the fit of the parametric models was assessed according to the discrepancy of the parametric (tail) density functions $g_1, \ldots, g_M$ defined over $(\boldsymbol{u}, \infty)$ and the histogram $\tilde{f}_{\boldsymbol{X}^{[u]}}$. Their tail index (generalised here to d dimensions) is given by

$$\tilde{T}_1(g_j) = \int_{(\boldsymbol{u},\infty)} |g_j(\boldsymbol{x}) - \tilde{f}_{\boldsymbol{X}^{[u]}}(\boldsymbol{x})| \mathrm{d}\boldsymbol{x}, \tag{4}$$

with the preferred models being those which give the smaller or smallest discrepancy

$$\underset{j \in \mathcal{M}}{\operatorname{argmin}}\ \tilde{T}_1(g_j).$$

Note that the subscript of $\tilde{T}_1$ indicates the L_1 error measure used in its definition (we will use L_2 and χ^2 error versions of this index later).

An obvious improvement to this procedure is to replace the histogram in Equation (4) with the transformation kernel estimator $\hat{f}_{\boldsymbol{X}^{[u]}}$. This will accordingly allow the usual the artefacts of histogram estimators to be avoided or at least reduced. These include the anchor point problem (i.e. how to specify the locations of the histogram bins) and the empty bin problem (where it is unclear whether histogram bins with empty counts should be interpreted as a true zero probability or are due to insufficient observed data). This latter case is important for extreme values as they are sparsely distributed in the tail regions. Specifically, the tail index would be calculated as

$$\hat{T}_1(g_j) = \int_{(\boldsymbol{u},\infty)} |g_j(\boldsymbol{x}) - \hat{f}_{\boldsymbol{X}^{[u]}}(\boldsymbol{x})| \mathrm{d}\boldsymbol{x}. \tag{5}$$

In the following section, we highlight the purpose of working with transformed density estimators, by directly contrasting $\tilde{T}_1(g_j)$ and $\hat{T}_1(g_j)$ with the index based on the standard kernel density estimator

$$\hat{T}_1^*(g_j) = \int_{(\boldsymbol{u},\infty)} |g_j(\boldsymbol{x}) - \hat{f}^*_{\boldsymbol{X}^{[\boldsymbol{u}]}}(\boldsymbol{x})| \mathrm{d}\boldsymbol{x}, \tag{6}$$

where $\hat{f}^*_{\boldsymbol{X}^{[\boldsymbol{u}]}}(\boldsymbol{x})$ represents the standard kernel density estimator constructed using the extreme data $\boldsymbol{X}_1^{[\boldsymbol{u}]}, \ldots, \boldsymbol{X}_n^{[\boldsymbol{u}]}$ only. The integrals in Equations (4)–(6) can be approximated by (weighted) Reimann sums.

3 Numerical studies

3.1 Simulated data - univariate

We now numerically examine the performance of the kernel density estimator introduced in Section 2.1 for univariate extremes, and demonstrate that it is a good surrogate for the true tail distribution. We additionally evaluate the estimator's performance through the model assessment procedure of Section 2.5. Gaussian kernels are adopted throughout given the usual secondary level of importance given to kernel choice in standard kernel methods (see e.g. Table 2.1 of Wand and Jones (1995) where the difference between the most efficient (Epanechnikov) kernel and the least efficient (uniform) kernel is less than 7%). However, kernel choice may be more important for extreme values, and so we return to this point below.

We generate a dataset of size $m = 2,000$ from each of the Gumbel, Fréchet and generalised Pareto (GPD) target distributions, and set the threshold u at the 95% upper sample quantile to produce a tail sample of size $n = 100$. For each sample, the target distribution is known exactly as $f_{X^{[u]}}(x) = f_X(x)/\bar{F}_X(u)$. For each sample, we compute:

1) The appropriate maximum-likelihood based, parametric estimator: Fréchet (FRE), Gumbel (GUM) or generalised Pareto (GPD);

2) The histogram estimator $\tilde{f}_{X^{[u]}}$ with normal scale optimal binwidth (HIS);

3) The transformation kernel based estimator $\hat{f}_{X^{[u]}}$ with transformation $t(x) = \log(x - u)$, using normal scale (KNS), plug-in (KPI), unbiased cross validation (KUC) and smoothed cross validation (KSC) optimal bandwidth selectors;

4) The standard kernel based estimators $\hat{f}^*_{X^{[u]}}$ based on the tail sample, using using normal scale (KNS*), plug-in (KPI*), unbiased cross validation (KUC*) and smoothed cross validation (KSC*) optimal bandwidth selectors.

The top row of Figure 1 illustrates the various tail density estimates for the three target distributions, with the true density shown as a solid black line. Displayed are the histogram estimator $\tilde{f}_{X^{[u]}}$ (HIS; red dashed line), and the transformed $\hat{f}_{X^{[u]}}$ (KPI; green dotted) and standard $\hat{f}^*_{X^{[u]}}$ (KPI*; blue dot-dash) kernel density estimates, both with plug in estimators only for clarity. Visually, the transformed kernel estimators appear to be the more accurate estimators of tail behaviour in each case. There is less obvious boundary bias, in particular compared to the standard kernel density estimator. The optimal histogram binwidth tends to be too small for these tail data, and as a result produces histograms which are more noisy compared to the kernel estimators. That the

kernel density estimators are naturally continuous kernel estimators also lead to better visualisation, and they are more helpful when comparing to a continuous target density.

The bottom row of Figure 1 examines the extremal performance of the same estimators through qq-plots of the target quantiles versus the non-parametric estimated quantiles. In all cases, the histogram estimator most consistently approximates the true quantiles. This performance compared to the kernel-based estimators is not unexpected, however, as the latter aim to optimally estimate the density function rather than the quantile function. Comparing the two kernel-based estimators, the transformed kernel estimator tends to either outperform (centre, right panels) or perform as well as (left panel) the standard estimator, which can be attributed to the standard estimator's natural boundary bias.

Note that the transformed kernel estimator has produced estimates with lighter tails than the true density. For large n this is possibly due to the choice of a Gaussian kernel K_h to construct the density estimates, so that the upper tail of this estimate (mapped through the inverse transform $\boldsymbol{t}^{-1}$) is light compared to the true Fréchet and Pareto tails. For smaller n, finite sample variation can produce a density estimate with either lighter or heavier tails in the body of the data (see Supplementary Information).

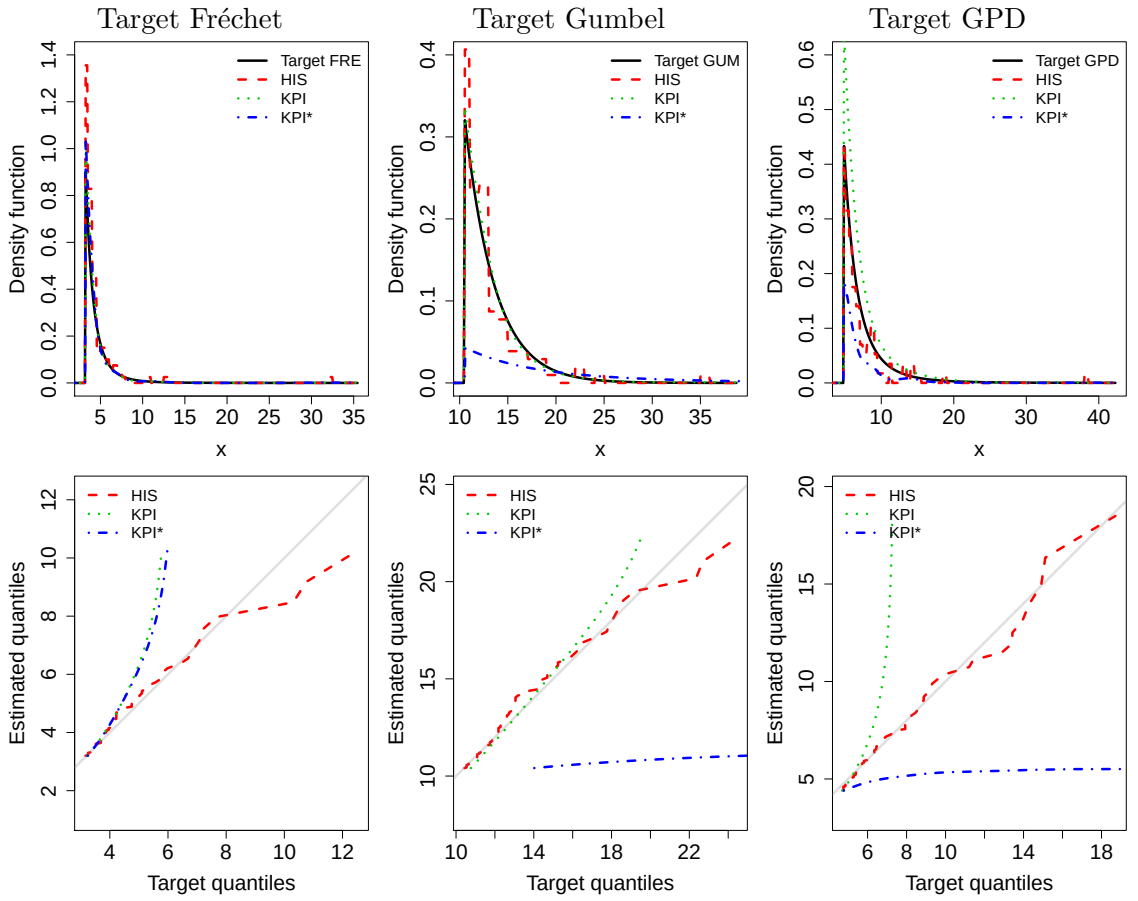

Figure 1: Non-parametric estimators of the univariate tail density (top) and of the tail quantiles (bottom) when the target density is Fréchet (left), Gumbel (centre) and generalised Pareto (right). Sample size is $n = 100$. Fréchet ($\mu = 1$, $\sigma = 0.5$, $\xi = 0.25$), Gumbel ($\mu = 1.5$, $\sigma = 3$) and Pareto ($\mu = 0$, $\sigma = 1$, $\xi = 0.25$) target densities are represented by a solid black line. The histogram estimator $\tilde{f}_{X^{[u]}}$ with normal scale binwidth (HIS) is represented by a dashed red line, the transformed kernel plug-in estimator $\hat{f}_{X^{[u]}}$ (KPI) by a dotted green line and the standard kernel estimator $\hat{f}^*_{X^{[u]}}$ (KPI*) by a dot-dash blue line.

For more quantitative results we repeat this process over 400 replicates for different dataset sizes $m = 500, 1000$ and 2000, producing tail samples of size $n = 25, 50$ and 100, with the threshold u set at the 95% upper quantile. As these three sample sizes gave similar results, we only present those for $n = 100$ for brevity. (See the Supplementary Information for results with $n = 25, 50$.) We take a numerical approximation (Reimann sum) of the L_2 error $\int_{(\boldsymbol{u},\infty)}[\text{estimator of} f_{\boldsymbol{X}^{[u]}}(\boldsymbol{x}) - f_{\boldsymbol{X}^{[u]}}(\boldsymbol{x})]^2 \, \mathrm{d}\boldsymbol{x}$, and of the χ^2 error $\int_{(\boldsymbol{u},\infty)}(\text{estimator of} f_{\boldsymbol{X}^{[u]}}(\boldsymbol{x}) - f_{\boldsymbol{X}^{[u]}}(\boldsymbol{x}))^2 / f_{\boldsymbol{X}^{[u]}}(\boldsymbol{x}) \mathrm{d}\boldsymbol{x}$ as measures of accuracy of each density estimator.

Figure 2 presents box-plots of the accuracy of each tail density estimator for each true tail distribution, and for both L_2 (top panels) and χ^2 (bottom) accuracy measures. As expected, for each target distribution the most accurate density estimator is the correctly specified parametric model. The transformed kernel density estimators systematically perform better than their standard kernel counterparts, although they can be more variable. The standard kernel estimators and the histogram estimator compete for the worst estimate of the tail density, depending on the true target distribution and the measure of estimator accuracy. The differences in the accuracy between kernel estimators with different bandwidth selectors is small, in contrast to studies where the bandwidth selection class is a crucial factor (see e.g. Sheather and Jones, 1991, Wand and Jones, 1995, Chapter 3), indicating that it is the difference between estimators that is dominating performance.

Finally, we examine density estimator performance in terms of it's ability to correctly select the true, data-generating model (Section 2.5). For each of the 3×400 datasets generated previously, we compute the tail indices $\tilde{T}_1$ (histogram estimator), $\hat{T}_1$ (transformed kernel), and $\hat{T}_1^*$ (standard kernel), modified to follow the form of L_2 loss (denoted by $\tilde{T}_2$, $\hat{T}_2$ and $\hat{T}_2^*$) and χ^2 loss ($\tilde{T}_{\chi^2}$, $\hat{T}_{\chi^2}$ and $\hat{T}_{\chi^2}^*$), with respect to each parametric model.

Table 1 displays the proportion of times that samples from a given true distribution are identified as coming from either Fréchet, Gumbel or GPD distributions (i.e. by having the smallest tail index value), as a function of tail density estimator. The proportion of replicates where the target model is correctly selected is given in bold. As the Gumbel distribution ($\xi = 0$) is on the limiting border of the parameter space of the Fréchet distribution ($\xi > 0$), to avoid possible model misidentification, we additionally perform a deviance test. If the Gumbel provides a significantly better fit than the Fréchet distribution, meaning that the shape parameter is not significantly different from zero, then we only consider Gumbel and GPD distributions as candidate models.

As might be expected, for any target distribution and error measure in Table 1, using kernel based estimators ($\hat{T}_\cdot$ and $\hat{T}_\cdot^*$) as a surrogate for the target density generally selects the correct target in the vast majority of cases, with proportions higher than those achieved through the histogram tail index $\tilde{T}_\cdot$. The one exception to this is for GPD data where the standard kernel estimator performs as poorly as the histogram estimator. The histogram-based tail index $\tilde{T}_2$ performs particularly poorly when the target density is Fréchet under L_2 error. This model selection inaccuracy is corrected when considering the χ^2 error, though accuracy is then lost with the GPD target. Overall, the transformation kernel-based index $\hat{T}_\cdot$ performs as strongly as, and in some cases much better than the standard kernel-based index $\hat{T}_\cdot^*$.

3.2 Simulated data - multivariate

The analysis of multivariate extremes is considerably more challenging than its univariate counterpart. A powerful motivation for exploratory data analysis using kernel-based estimation is that no single parametric family exists for max-stable distributions. See e.g.

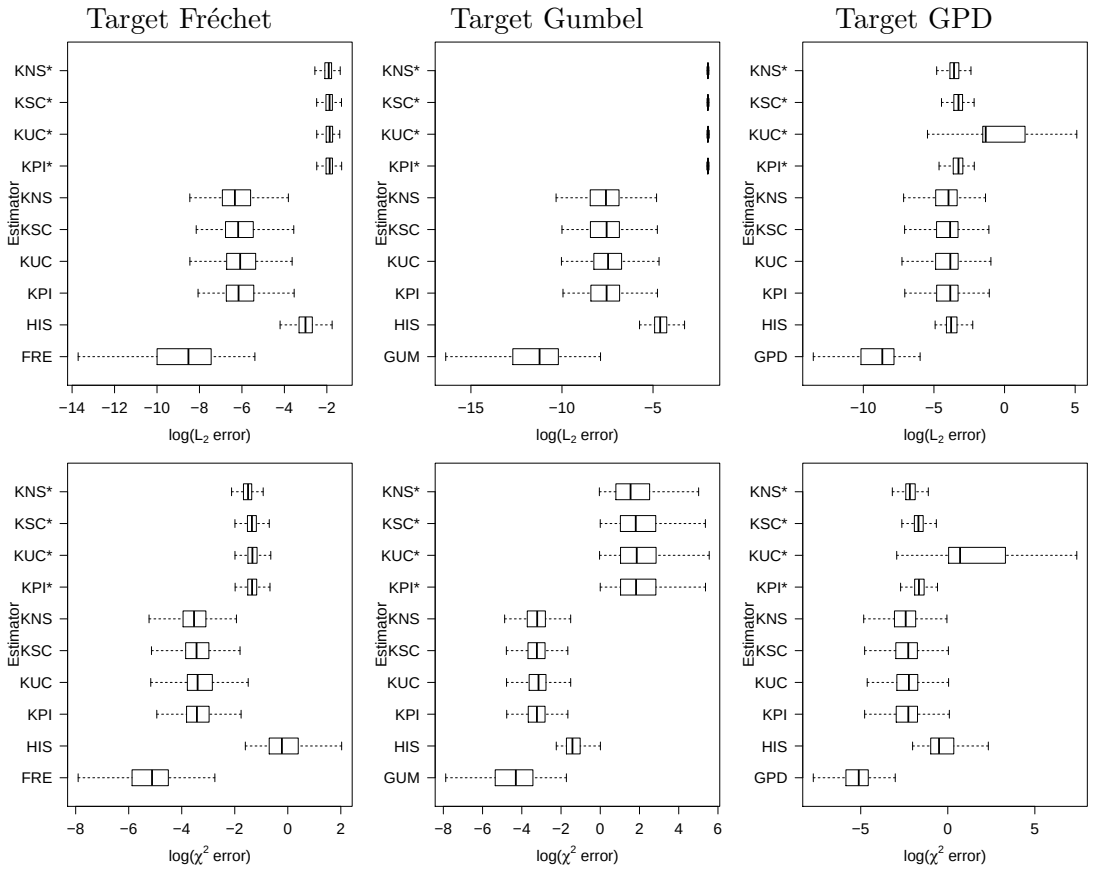

Figure 2: Box-plots of the $\log L_2$ (top panels) and $\log \chi^2$ (bottom) errors for the parametric Fréchet (FRE), Gumbel (GUM), generalised Pareto (GPD), and histogram (HIS) tail density estimators. Transformed kernel density estimators $\hat{f}_{X^{[u]}}$ use the plug-in (KPI), unbiased cross validation (KUC), smoothed cross validation (KSC) and normal scale kernel (KNS) optimal bandwidth selectors. Standard kernel density estimators $\hat{f}^*_{X^{[u]}}$ are indicated by an asterisk (*). True target densities are (left panels) Fréchet, (centre) Gumbel and (right) GPD. Box plots are based on 400 replicates of $m = 2,000$ observations with a tail sample size of $n = 100$.

		$\tilde{T}_\cdot$ selected model			$\hat{T}_\cdot$ selected model			$\hat{T}^*_\cdot$ selected model		
Error	Target	FRE	GUM	GPD	FRE	GUM	GPD	FRE	GUM	GPD
L_2	FRE	**0.03**	0.18	0.79	**0.92**	0.00	0.08	**0.89**	0.00	0.11
	GUM	0.02	**0.90**	0.08	0.04	**0.94**	0.02	0.02	**0.98**	0.00
	GPD	0.01	0.28	**0.71**	0.11	0.05	**0.84**	0.00	0.31	**0.69**
χ^2	FRE	**1.00**	0.00	0.00	**1.00**	0.00	0.00	**1.00**	0.00	0.00
	GUM	0.04	**0.94**	0.02	0.04	**0.96**	0.00	0.04	**0.96**	0.00
	GPD	0.78	0.00	**0.22**	0.20	0.00	**0.80**	0.58	0.00	**0.42**

Table 1: Proportion of 400 simulated datasets from each known target distribution (Fréchet, Gumbel and GPD) that are identified as coming from each of these distributions by having the smallest tail index value, as a function of nonparametric density estimator. Bold text indicates the proportion selecting the correct model. Nonparametric density estimators are the histogram ($\tilde{T}_\cdot$), the transformed kernel ($\hat{T}_\cdot$) and the standard kernel ($\hat{T}^*_\cdot$), with tail indices calculated according to the L_2 (top row) and χ^2 loss (bottom row).

Kotz and Nadarajah (2000), Coles (2001), Beirlant et al. (2004), de Haan and Ferreira (2006), Falk et al. (2011) and Beranger and Padoan (2015) for theoretical details and applications.

We now numerically examine the performance of a bivariate transformation-based density estimator, similarly to that for the univariate case. We generate datasets of size $m = 4,000$ from the asymmetric negative logistic (ANL; Joe, 1990), the bilogistic (BIL; Smith et al., 1990) and the Hüsler-Reiss (HR; Hüsler and Reiss, 1989) distributions. The bivariate extreme sample is then determined from a $n = 200$ random selection of the observations that lie within $(\boldsymbol{u}, \infty)$, where $\boldsymbol{u}$ is constructed from the dataset's marginal 90% upper quantiles.

For each sample we compute the appropriate maximum likelihood based parametric estimator (assuming simultaneously estimated generalised extreme value distribution margins), a 2-dimensional histogram with normal scale optimal bandwidth (HIS) and the transformation and standard kernel estimators with plug-in optimal bandwidth selectors (respectively KPI and KPI*) and transformation $\boldsymbol{t}(\boldsymbol{x}) = (\log(x_1 - u_1), \log(x_2 - u_2))^\top$. The results from the other bandwidth selectors are not displayed both for clarity, and due to the limited impact of the bandwidth selector method on the performance of the density estimator.

Figure 3 illustrates the 25%, 50%, 75% and 99% highest density level sets of the histogram (red dashed line), transformed kernel density estimates (dotted green line) and standard kernel density estimates (dot-dashed blue line) in comparison with the know target distribution (solid black line). Visually, the transformed kernel estimator performs well – it is able to identify and describe most of the features of the target densities, although it doesn't quite capture the slight asymmetry for the asymmetric negative logistic target. In contrast, the blocky, discrete nature of the histogram estimator makes it difficult to discern the nature of the underlying target, and the standard kernel estimator is clearly unable to capture the features of the tail density as accurately as the transformation kernel estimator.

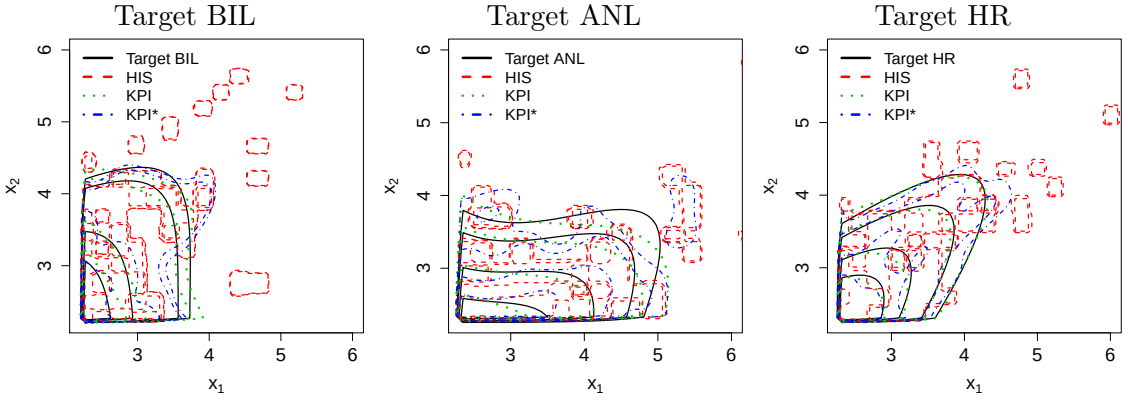

Figure 3: Non-parametric estimators of the bivariate tail density when the target density is bilogistic (BIL), asymmetric negative logisitic (ANL) and Hüsler-Reiss (HR). Sample size is $n = 200$. Bilogistic ($\alpha = 0.8$, $\beta = 0.52$), asymmetric negative logistic (dependence parameter = 1.3, asymmetry parameter = $(0.2, 0.7)$) and Hüsler-Reiss (dependence parameter = 2.4) target quantiles are represented by a solid black line. The histogram estimator $\tilde{f}_{\boldsymbol{X}^{[u]}}$ with a normal scale bin width (HIS) is represented by a dashed red line, the transformation kernel estimator $\hat{f}_{\boldsymbol{X}^{[u]}}$ with plug-in bandwidth estimator (KPI) by the dotted green line and the standard kernel estimator $\hat{f}^*_{\boldsymbol{X}^{[u]}}$ with plug-in bandwidth estimator (KPI*) by the dot-dashed blue line.

Figure 4 measures the $\log L_2$ and $\log \chi^2$ performance of the parametric (BIL, ANL,

HR), histogram (HIS), and transformation and standard kernel (KPI and KPI*) tail density estimators in approximating the known target distribution, for each of the target distributions considered in Figure 3. As for the univariate case, the correctly specified parametric estimator of each distribution generates the smallest error. The transformation kernel density estimator produces the next most efficient estimator, with the bivariate histogram and the bivariate standard kernel performing the most poorly in each case. The differences between the non-parametric estimators is the smallest for the asymmetric negative logistic model.

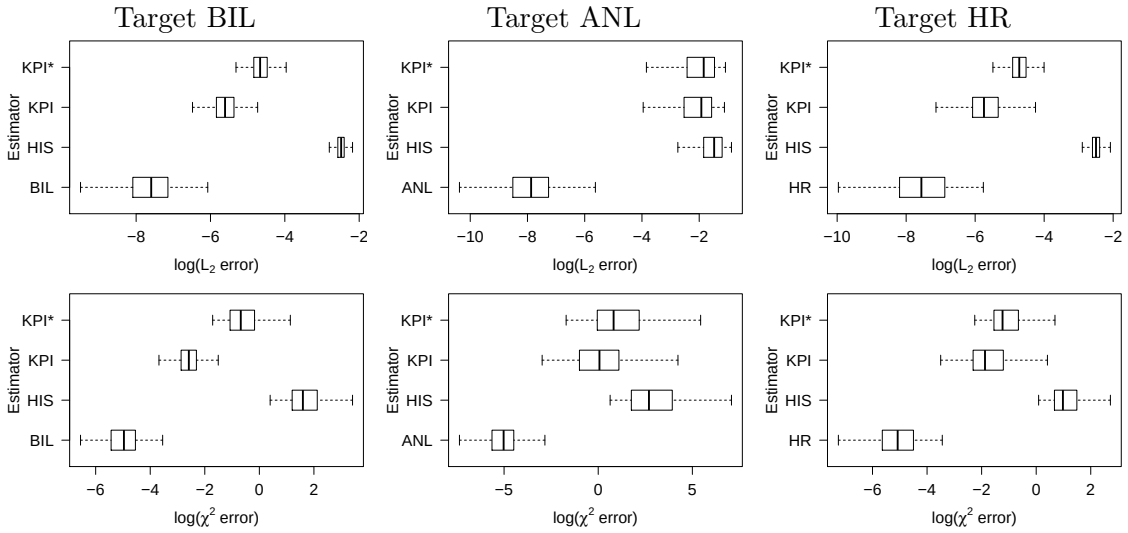

Figure 4: Box-plots of the $\log L_2$ (top panels) and $\log \chi^2$ (bottom) errors for the bivariate biologistic (BIL), asymmetric negative logistic (ANL) and Hüsler-Reiss (HR) parametric estimators, the 2-dimensional histogram (HIS), and the bivariate transformation kernel estimator with plug-in optimal bandwidth selector (KPI). True target densities are (left panels) the bivariate biologistic, (centre) the asymmetric negative logistic and (right) the Hüsler-Reiss models. Box plots are based on 400 replicates of $m = 4,000$ observations with a tail sample size of $n = 200$.

Finally, we examine the ability of the density estimator to correctly select the true data generating model. Similarly to Table 1, Table 2 shows the proportion of time that each model was selected based on the bivariate histogram ($\tilde{T}_\cdot$), transformation kernel ($\hat{T}_\cdot$) and standard kernel $\hat{T}_\cdot^*$ based tail indices using data generated from a known model. The bold figures indicate the proportion of replicates selecting the correct model.

Overall, the results are mixed. Unsurprisingly, the standard kernel based index is able to choose the correct model more consistently than the histogram based index. The exception to this is for the Hüsler-Reiss target under χ^2 error where the standard kernel estimator performs comparatively poorly. Between the two kernel-based estimators, the best performer changes depending on the target distribution and error measure. The transformation density estimator uniformly and most consistently selects the true data generating model for Hüsler-Reiss data (where the standard kernel estimator performed quite poorly), and for ANL data under χ^2 error. However, it underperforms for biologistic and ANL generated data under L_2 error, although all methods have difficulty identifying the ANL model in general.

4 Exploratory data analysis of climate models

Perkins et al. (2007, 2013) previously used univariate histogram density estimators for both

		$\tilde{T}_\cdot$ selected model			$\hat{T}_\cdot$ selected model			$\hat{T}_\cdot^*$ selected model		
Error	Target	BIL	ANL	HR	BIL	ANL	HR	BIL	ANL	HR
L_2	BIL	**0.81**	0.06	0.13	**0.68**	0.06	0.26	**0.84**	0.04	0.12
	ANL	0.84	**0.16**	0.00	0.97	**0.03**	0.00	0.85	**0.15**	0.00
	HR	0.24	0.03	**0.73**	0.00	0.05	**0.95**	0.16	0.02	**0.82**
χ^2	BIL	**0.66**	0.04	0.30	**0.66**	0.06	0.28	**0.70**	0.23	0.07
	ANL	0.72	**0.21**	0.07	0.52	**0.48**	0.00	0.56	**0.43**	0.01
	HR	0.00	0.02	**0.98**	0.00	0.00	**1.00**	0.09	0.27	**0.64**

Table 2: Proportion of 400 simulated datasets from each known target distribution (BIL bilogistic, ANL asymmetric negative logistic, and HR Hüsler-Reiss) that are identified as coming from each of these distributions by having the smallest tail index value, as a function of nonparametric density estimator. Bold text indicates the proportion selecting the correct model. Nonparametric density estimators are the bivariate histogram ($\tilde{T}_\cdot$), the transformed kernel ($\hat{T}_\cdot$) and the standard kernel ($\hat{T}_\cdot^*$), with tail indices calculated according to the L_2 (top row) and χ^2 (bottom row) loss.

visualisation and model selection to identify strongly and poorly performing global climate models (GCMs) in terms of their ability to reproduce extreme (minimum and maximum) temperatures in Australia. The models which they considered are the climate models assessed by the Intergovernmental Panel on Climate Change (IPCC) Fourth Assessment Report (AR4) to investigate changes in temperature extremes. A well-known challenge for these models is to be able to accurately project extreme temperatures (Perkins et al., 2007, 2013, Sillmann et al., 2013a,b, Cowan et al., 2014, Fischer et al., 2013).

Following earlier work, Perkins et al. (2013) developed a univariate tail index (see Section 2.5, equation (4)) which evaluates the amount of overlap between a model-predicted distributional tail, g_i, and the distribution of the extreme observed data. This index reflects the discrepancy between two distributional tails, whereby a model perfectly fitting the observed data has zero score, and increasing scores imply an increasing lack-of-fit of the model to the observed data. Unlike for the simulated parametric models in Section 3.1 and 3.2 , there is no closed form for the density function $g_i, i = 1, \ldots, M$, to characterise the data values generated by the climate models. Perkins et al. (2013) replaced the unknown target density g_i with a histogram $\tilde{g}_i$, based on model generated data, when comparing to the histogram of the observed data $\tilde{f}_{X^{[u]}}$ in Equation (4). I.e. they used the index $\tilde{T}_1(\tilde{g}_i)$ to determine the most appropriate model

We extend this analysis by considering a wider and more modern ensemble of global climate models than those in Perkins et al. (2013), as well as exploring alternatives to $\tilde{T}_1(\tilde{g}_i)$ as the model selection criterion. Here we use $M = 22$ climate models participating in the World Climate Research Programme's 5th phase Coupled Model Intercomparison Project (CIMP5; see Flato et al., 2013), which currently underpin global and regional climate projections of extremes (e.g. Sillmann et al., 2013b). The choice of models was based on the availability of daily maximum and minimum temperature data for the historical experiment ($\sim$ 1860-2005 ; see Taylor et al., 2012). Other targeted temperature extreme evaluation studies on the CMIP5 ensemble have found generally well-simulated changes in observed trends of specific indices (e.g. Sillmann et al., 2013a, Flato et al., 2013). However unlike this study, no consideration has been given to the full underlying distribution of extremes.

The observed data sample are the daily observed maximum temperatures observed in the Sydney region in Australia from 01/01/1911 to 31/12/2005 yielding a sample of $m = 34,699$ observations. Observations were obtained from the Australian Water Availability Project dataset (Jones et al., 2009), a gridded product covering all of Australia. All AR4

climate models were run to generate data in this same time frame, and the GCM grid box in which Sydney is located was extracted. The threshold determining the extreme maximum temperatures is the 95% upper quantile $u = 30.98°$C, producing a tail sample of size $n = 1,727$.

Table 3 displays the modified Perkins et al. (2013) histogram-based tail indices, $\tilde{T}_2(\tilde{g}_i)$ and $\tilde{T}_{\chi^2}(\tilde{g}_i)$, and the transformation kernel density estimator based equivalents, $\hat{T}_2(\hat{g}_i)$ and $\hat{T}_{\chi^2}(\hat{g}_i)$, for each of the $i = 1, \ldots, 22$ climate models. Note that $\hat{T}_\cdot(\hat{g}_i)$ implements the transformation kernel estimator for both the observed data ($\hat{f}_{X^{[u]}}$) and the GCM generated data (within $\hat{T}_\cdot$). The bold values in each column indicate the three smallest index scores.

From Table 3, both histogram- and kernel-based indices strongly identify MPI-ESM-LR and MPI-ESM-MR as the top two performing models under both L_2 and χ^2 error. The third model clearly identified by $\tilde{T}_\cdot(\tilde{g})$ is CMCC-CESM, whereas the HadGEM2-ES and CNRM-CMS models are respectively selected by $\hat{T}_2(\hat{g})$ and $\hat{T}_{\chi^2}(\hat{g})$. It is useful to note that for all indices, the ACESS1-3, NorESM1-M, IPSL-CM5* and MIROC-ESM* models produce inadequate fits.

	Model selection index			
Model	$\tilde{T}_2(\tilde{g})$	$\tilde{T}_{\chi^2}(\tilde{g})$	$\hat{T}_2(\hat{g})$	$\hat{T}_{\chi^2}(\hat{g})$
ACCESS1-0	0.0013	0.1057	0.0008	0.0646
ACCESS1-3	0.0089	2.0474	0.0079	6.3443
CanESM2	0.0007	0.0812	0.0002	0.0152
CMCC-CESM	**0.0005**	**0.0648**	0.0002	0.0265
CMCC-CM	**0.0006**	0.0665	0.0002	0.0190
CMCC-CMS	0.0010	0.0964	0.0005	0.0462
CNRM-CMS	0.0011	0.0814	0.0002	**0.0146**
CSIRO-Mk3-6-0	0.0008	0.0833	0.0005	0.0398
HadCM3	0.0008	0.0953	0.0005	0.0687
HadGEM2-CC	0.0006	0.0856	0.0002	0.0278
HadGEM2-ES	0.0007	0.0758	**0.0001**	0.0172
IPSL-CM5A-LR	0.0056	0.2740	0.0037	0.1639
IPSL-CM5A-MR	0.0023	0.1299	0.0011	0.0556
IPSL-CM5B-LR	0.0046	0.2429	0.0032	0.1629
MIROC-ESM	0.0081	0.4835	0.0059	0.3028
MIROC-ESM-ECHM	0.0074	0.4581	0.0051	0.2830
MIROC5	0.0010	0.0779	0.0003	0.0188
MPI-ESM-LR	0.0007	**0.0583**	**0.0002**	**0.0052**
MPI-ESM-MR	**0.0005**	**0.0568**	**0.0001**	**0.0051**
MPI-ESM-P	0.0012	0.1045	0.0008	0.0440
MRI-CGCM3	0.0009	0.1209	0.0007	0.1128
NorESM1-M	0.0039	0.5759	0.0008	0.1381

Table 3: Univariate histogram- $\tilde{T}_\cdot(\tilde{g}_i)$ and kernel-based $\hat{T}_\cdot(\hat{g}_i)$ tail index scores, based on histogram $\tilde{g}_i$ and kernel $\hat{g}_i$ density estimators, for the extreme maximum temperatures produced by the $i = 1, \ldots, 24$ AR4 climate models. Both L_2 and χ^2 error measures are implemented. Bold values in each column indicate the top three models with smallest tail index values.

Figure 5 illustrates both the tail density estimators and qq-plots for four of the top performing models. The histogram/kernel estimates are represented by the solid/dashed lines, whereas the observed and GCM data are displayed with black and grey lines respectively. For the top two models (MPI-ESM-LR and MPI-ESM-MR), the histogram

and kernel density estimator of the simulated data closely follow the histogram and kernel density estimator of the observed data. This is supported by the qq-plots. While the MPI-ESM-LR and MPI-ESM-MR climate models closely fit the observed data, the CMCC-CESM and HadGEM2-ES models visibly fit the observed data less well. While for any fixed model, it is the differences between the climate model output and observed data rather than form of the tail density estimator that is dominating the choice of most accurate model, visually, the HadGEM2-ES (selected by $\hat{T}_2(\hat{g})$) model appears to be a better fit to the observed data than CMCC-CESM (selected by $\tilde{T}_\cdot(\tilde{g})$), which may arguably lend weight to the kernel based index.

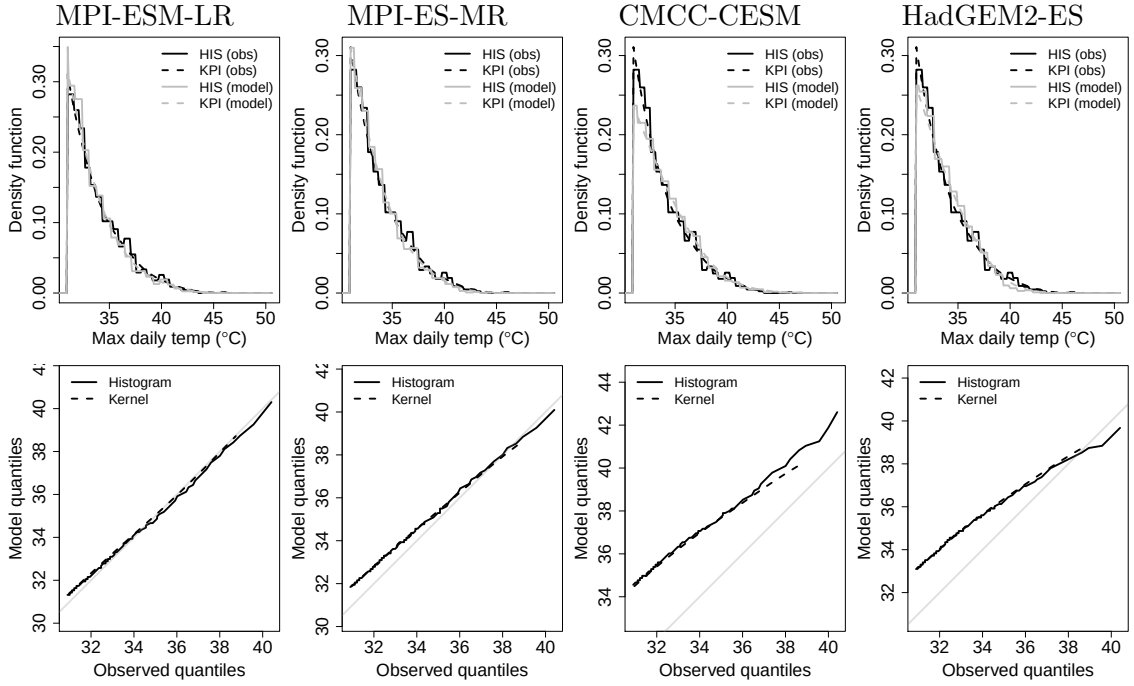

Figure 5: Histogram ($\tilde{g}_i$) and transformation kernel ($\hat{f}_{X^{[u]}}$) based estimators of the tail densities for four top AR4 models: (left to right) MPI-ESM-LR, MPI-ES-MR, CMCC-CESM and HadGEM2-ES. Histogram estimators (HIS) are denoted by solid lines and kernel plug-in estimators (KPI) by dashed lines. Observed data (obs) is illustrated in black and GCM data (model) in grey.

For a bivariate analysis, we consider the upper tail of pairs of maximum and minimum temperatures over the same time period, in order to investigate which of the climate models can predict a joint increase. (The largest minimum temperatures are important in understanding the duration and severity of heatwaves.) The threshold for the maximum temperatures are the 90% marginal upper quantiles $\boldsymbol{u} = (28.77°\text{C}, 18.07°\text{C})^\top$ with a tail sample size of $n = 1,387$. Table 4 presents the same information as Table 3 but for the bivariate data.

The CNRM-CMS model is identified by all four indices and achieved 3 out of 4 of the lowest tail index scores. HadCM3 and NorESM1-M are also identified by both $\tilde{T}_\cdot(\hat{g})$ and $\hat{T}_\cdot(\hat{g})$ (the former by $\hat{T}_2(\hat{g})$, the latter by $\hat{T}_{\chi^2}(\hat{g})$) as competitive models for the prediction of extreme minimum and maximum temperatures. Note these are different models to those that were selected for the univariate analysis. Indeed, the NorESM1-M model which performed quite poorly for the marginal univariate maximum temperature analysis, has now becomes a viable model when considering the joint (maximum, minimum) extremes. This likely points to a large variation in the accuracy of GCM predictions across the range

of GCM data outputs.

	Model selection index			
Model	$\tilde{T}_2(\tilde{g})$	$\tilde{T}_{\chi^2}(\tilde{g})$	$\hat{T}_2(\hat{g})$	$\hat{T}_{\chi^2}(\hat{g})$
ACCESS1-0	4.8620e^{-2}	4.1304e^{0}	1.7793e^{-3}	3.0529e^{0}
ACCESS1-3	2.8232e^{-2}	2.7560e^{0}	3.6377e^{-3}	1.1024e^{2}
CanESM2	2.2005e^{-2}	1.9621e^{0}	1.7774e^{-3}	3.8483e^{0}
CMCC-CESM	7.5085e^{-2}	5.7584e^{0}	2.1369e^{-3}	4.4225e^{0}
CMCC-CM	6.6515e^{-3}	5.2560e^{-1}	$\mathbf{8.3709}\text{e}^{\mathbf{-4}}$	4.4277e^{-1}
CMCC-CMS	1.4774e^{-1}	1.1375e^{1}	4.1119e^{-3}	1.1520e^{1}
CNRM-CMS	$\mathbf{3.4427}\text{e}^{\mathbf{-3}}$	$\mathbf{1.9220}\text{e}^{\mathbf{1}}$	$\mathbf{4.6494}\text{e}^{\mathbf{-4}}$	$\mathbf{1.8326}\text{e}^{\mathbf{-1}}$
CSIRO-Mk3-6-0	4.8413e^{-2}	8.4685e^{-1}	1.8034e^{-3}	3.4450e^{0}
HadCM3	$\mathbf{4.1697}\text{e}^{\mathbf{-3}}$	$\mathbf{0.3499}\text{e}^{\mathbf{0}}$	$\mathbf{7.2225}\text{e}^{\mathbf{-4}}$	7.9124e^{-1}
HadGEM2-CC	3.5372e^{-2}	2.9277e^{0}	1.5636e^{-4}	2.0376e^{0}
HadGEM2-ES	3.3287e^{-2}	2.7451e^{0}	1.2882e^{-4}	1.9663e^{0}
IPSL-CM5A-LR	1.1441e^{-2}	1.0183e^{0}	1.4796e^{-4}	2.7279e^{0}
IPSL-CM5A-MR	1.7574e^{-2}	1.4335e^{0}	1.4112e^{-4}	4.5790e^{0}
IPSL-CM5B-LR	1.6266e^{-2}	1.3757e^{0}	1.4201e^{-4}	$\mathbf{4.3748}\text{e}^{\mathbf{-1}}$
MIROC-ESM	2.5876e^{-2}	2.6036e^{0}	1.2148e^{-2}	9.6383e^{1}
MIROC-ESM-ECHM	3.3514e^{-2}	3.5886e^{0}	1.2829e^{-2}	1.0937e^{2}
MIROC5	4.8100e^{-2}	3.9027e^{0}	1.0442e^{-2}	6.7268e^{1}
MPI-ESM-LR	1.5531e^{-2}	1.2400e^{0}	1.7581e^{-3}	1.8988e^{0}
MPI-ESM-MR	3.1483e^{-2}	2.5861e^{0}	3.9513e^{-3}	6.6629e^{0}
MPI-ESM-P	1.7825e^{-2}	1.2888e^{0}	2.3292e^{-3}	2.3445e^{0}
MRI-CGCM3	5.5388e^{-2}	4.0101e^{0}	4.7491e^{-3}	2.9753e^{1}
NorESM1-M	$\mathbf{6.3489}\text{e}^{\mathbf{-3}}$	$\mathbf{3.0830}\text{e}^{\mathbf{-1}}$	1.1626e^{-3}	$\mathbf{1.1540}\text{e}^{\mathbf{-1}}$

Table 4: Bivariate histogram- $\tilde{T}(\tilde{g}_i)$ and kernel-based $\hat{T}(\hat{g}_i)$ tail index scores, based on histogram $\tilde{g}_i$ and kernel $\hat{g}_i$ density estimators, for the extreme (minimum, maximum) temperatures produced by the $i = 1, \ldots, 24$ AR4 climate models. Both L_2 and χ^2 error measures are implemented. Bold values in each column indicate the top three models with smallest tail index values.

Figure 6 illustrates the plots of the bivariate tail density estimators for the CNRM-CMS and NorESM1-M models. Similarly to Figure 5, histograms are shown by solid lines, and kernel estimators by dashed lines, with the observed data represented by black lines and the GCM data by grey. It is immediately apparent that the kernel-based density estimates are visually much cleaner, and easier to evaluate than their histogram counterparts.

Both GCM models appear to simulate the extremes of the maximum temperatures quite well, however for NorESM1-M the magnitude of the minimum temperatures is underestimated and the maximum temperatures are overestimated. This indicates two possibilities. That the kernel estimators need to be further refined at the boundary or, perhaps more likely, that physical parameters in the GCMs need to be revised for a more realistic simulation of minimum temperature extremes. While minimum temperatures are somewhat easier for a climate model to simulate than maximum temperature (Perkins et al., 2007), the mis-representation of the observed temperature distribution is a well-known issue for GCMs, which is at least in part explained by their coarse resolution (Seneviratne et al., 2012). For example, dynamically downscaled regional climate models that are run at finer resolutions for a limited spatial domain can offer some improvement in the simulation of extreme temperatures (e.g. Seneviratne et al., 2012, Vautard et al., 2013, Perkins

et al., 2014).

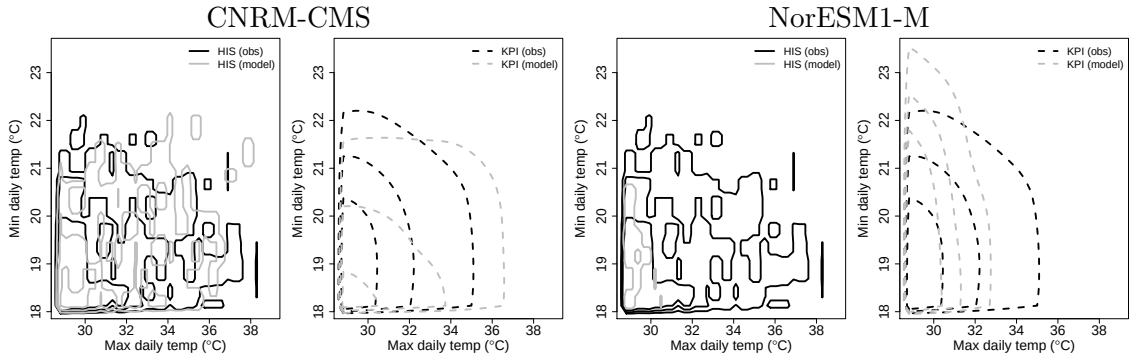

Figure 6: Histogram (left panel) and kernel estimators (right panel) of the tail densities for the CNRM-CMS and NorESM1-M climate models. Histograms estimators (HIS) are denoted by solid lines and kernel plug-in estimators (KPI) by dashed lines. Observed data (obs) is illustrated with black lines and climate model data (model) by grey lines. Kernel estimator contours indicate the 25%, 50% and 75% highest density level sets.

5 Discussion

Nonparametric density estimation is a useful exploratory data analysis tool for extremes. In this article we have introduced a non-parametric kernel estimator for the analysis of the tail density of univariate and multivariate data by applying a monotonic transformation $\boldsymbol{t}(\boldsymbol{x})$ to account for the boundedness and skewness of the tail sample on $(\boldsymbol{u}, \infty)$. Our theoretical results (centred on Theorem 1) indicate that the performance of this transformation kernel density estimator is comparable to that of standard kernel methods on non-truncated data in terms of minimal MISE, and that it outperforms standard kernel methods applied directly to the tail-truncated data. Numerical illustrations of its performance were given in Sections 3 and 4.

This tail density estimator provides visually very useful representations of extreme sample behaviour compared to, say, histogram estimators – consider the contrast in visual clarity between the histogram and kernel estimators illustrated in Figure 6. As a result it can reliably be incorporated into existing diagnostic and performance measures, such as the tail index of Perkins et al. (2013).

There is, of course, scope for further development and analysis of these ideas. Throughout we have constructed our kernel density estimates based on Gaussian kernels, $K_{\mathbf{H}}$. In general, this means that the extreme tail behaviour of the kernel-based density estimators is necessarily the tail behaviour of $K_{\mathbf{H}}$ mapped through $\boldsymbol{t}^{-1}$ the inverse of the transformation $\boldsymbol{t}$. This implies that if the tail behaviour of the sample does not correspond to that of the transformed kernel, then the kernel density estimate will poorly represent the true behaviour of the observed data distribution beyond the range of the observed data. A natural approach to resolving this problem could be to adapt the form of $K_{\mathbf{H}}$ and $\boldsymbol{t}$ to directly correspond to the (estimated) tail behaviour of the observed sample.

Additional improvements could be obtained by incorporating a local polynomial adjustment (e.g. Geenens, 2014) to the lower boundary, although here our primary interest is in the behaviour of the upper tail. Similarly, while the logarithm transformation is widely used due it conveniently mapping a semi-infinite interval to the real line, alternative transformations could be considered (though note that the result of Theorem 1 does

not depend on the form of $\boldsymbol{t}(\boldsymbol{x})$). Possibilities include the shifted power family of Wand et al. (1991)

$$t(x) = \begin{cases} (x + \lambda_1)^{\lambda_2} \operatorname{sign}(\lambda_2) & \lambda_2 \neq 0 \\ \log(x + \lambda_2) & \lambda_2 = 0 \end{cases}$$

where the log-transformation fixes $\lambda_1 = 0, \lambda_2 = -u$, and so other values may lead to better estimation, and the richer family of transformations proposed by Marron and Ruppert (1994) (although see also Geenens, 2014).

Finally, the performance of this transformation-based kernel density estimator is obviously limited by the performance of standard kernel density estimator methods. In particular, its performance will decline as the dimension of the random vector $\boldsymbol{X}$ increases. While this is unavoidable, if one wishes to perform kernel density estimation in this setting, it is important that it is implemented as efficiently as possible. The results presented in this article provide one step towards achieving this.

Acknowledgements

SEP and SAS are supported by the Australian Centre of Excellence for Climate System Science (CoECSS), SEP by the Australian Research Council grant number DE140100952, and SAS by the Australian Centre of Excellence for Mathematical and Statistical Frontiers in Big Data, Big Models, New Insights (ACEMS).

Appendix – Proofs

The below assumptions will be used to establish the optimality properties of our transformation kernel density estimators. These assumptions are usually expressed for random variables with unbounded support, which in our case is the transformed variable $\boldsymbol{Y} = \boldsymbol{t}(\boldsymbol{X}^{[\boldsymbol{u}]})$. This set of conditions do not form a minimal set, but they serve as a convenient starting point to state our results.

(A1) The d-variate density $f_{\boldsymbol{Y}}$ is continuous, square integrable and ultimately monotone for all element-wise partial second derivatives.

(A2) The d-variate kernel K is a positive, symmetric, square integrable p.d.f. such that $\int_{\mathbb{R}^d} \boldsymbol{y}\boldsymbol{y}^T K(\boldsymbol{y}) d\boldsymbol{y} = m_2(K)\mathbf{I}_d$ where $m_2(K)$ is finite and $\mathbf{I}_d$ is the $d \times d$ identity matrix.

(A3) The bandwidth matrix $\mathbf{H} = \mathbf{H}(n)$ forms a sequence of symmetric and positive definite matrices such that $n^{-1}|\mathbf{H}|^{-1/2}$ and every element of $\mathbf{H}$ approaches zero as $n \to \infty$.

Lemma 1 (minimal rate of MISE convergence)

The proof of Theorem 1 requires Lemma 1 (below) which establishes the minimal rate of MISE convergence of $\hat{f}_{\boldsymbol{Y}}$. This result has already been established (e.g. Wand, 1992), however we include details of a proof using an alternative notation for fourth order derivatives of a multivariate function via four-fold Kronecker product, which is simpler to code than tensors (see Holmquist, 1996).

Lemma 1. *Suppose that the conditions (A1–A3) hold. The minimal MISE of the transformed kernel density estimator* $\hat{f}_{\boldsymbol{Y}}(\cdot;\mathbf{H})$ *is*

$$\begin{aligned}\inf_{\mathbf{H}\in\mathcal{F}} \operatorname{MISE} \hat{f}_{\boldsymbol{Y}}(\cdot;\mathbf{H}) &= \left[\tfrac{1}{4}m_2^2(K)(\operatorname{vec}^T \mathbf{H}^*)^{\otimes 2}\boldsymbol{\psi}_{\boldsymbol{Y},4} + n^{-1}|\mathbf{H}^*|^{-1/2}R(K)\right]\{1+o(1)\}\\ &= O(n^{-4/(d+4)})\end{aligned}$$

where $\boldsymbol{\psi}_{\boldsymbol{Y},4} = \int_{\mathbb{R}^d} \mathsf{D}^{\otimes 4} f_{\boldsymbol{Y}}(\boldsymbol{y}) f_{\boldsymbol{Y}}(\boldsymbol{y}) \mathrm{d}\boldsymbol{y}$.

Proof of Lemma 1. The expected value of $\hat{f}_{\boldsymbol{Y}}$ is

$$\mathbb{E}\, \hat{f}_{\boldsymbol{Y}}(\boldsymbol{y};\mathbf{H}) = \mathbb{E}\, K_{\mathbf{H}}(\boldsymbol{y}-\boldsymbol{Y}) = \int_{\mathbb{R}^d} K_{\mathbf{H}}(\boldsymbol{y}-\boldsymbol{w}) f(\boldsymbol{y})\, d\boldsymbol{w} = K_{\mathbf{H}} * f_{\boldsymbol{Y}}(x)$$

where $*$ denotes the convolution operator between two functions. Asymptotically, using a Taylor series expansion and the usual rearrangement $\boldsymbol{w}^T\mathbf{H}^{1/2}\mathsf{D}^2F(\boldsymbol{x})\mathbf{H}^{1/2} = \operatorname{tr}(\boldsymbol{w}\boldsymbol{w}^T\mathbf{H}\mathsf{D}^2F(\boldsymbol{x}))$,

$$\begin{aligned}\mathbb{E}\, \hat{f}_{\boldsymbol{Y}}(\boldsymbol{y};\mathbf{H}) &= \int_{\mathbb{R}^d} |\mathbf{H}|^{-1/2} K(\mathbf{H}^{-1/2}(\boldsymbol{y}-\boldsymbol{w})) f_{\boldsymbol{Y}}(\boldsymbol{y}) d\boldsymbol{w}\\ &= \int_{\mathbb{R}^d} K(\boldsymbol{w}) f(\boldsymbol{y}-\mathbf{H}^{1/2}\boldsymbol{w}) d\boldsymbol{w}\\ &= \int_{\mathbb{R}^d} K(\boldsymbol{w})[f(\boldsymbol{y}) - \boldsymbol{w}^T\mathbf{H}^{1/2}\mathsf{D}f(\boldsymbol{x}) + \tfrac{1}{2}\boldsymbol{w}^T\mathbf{H}^{1/2}\mathsf{D}^2 f(\boldsymbol{y})\mathbf{H}^{1/2}\boldsymbol{w}]\{1+o(1)\}\, d\boldsymbol{w}\\ &= [f(\boldsymbol{y}) + \tfrac{1}{2}\int K(\boldsymbol{w}) \operatorname{tr}(\boldsymbol{w}\boldsymbol{w}^T\mathbf{H}\mathsf{D}^2 f(\boldsymbol{y}))\, d\boldsymbol{w}]\{1+o(1)\}\\ &= [f(\boldsymbol{y}) + \tfrac{1}{2}m_2(K)\operatorname{tr}(\mathbf{H}\mathsf{D}^2 f(\boldsymbol{y}))]\{1+o(1)\}.\end{aligned}$$

This allows us to write the bias of $\hat{f}_{\boldsymbol{Y}}(\boldsymbol{y};\mathbf{H})$ as

$$\mathbb{E}\, \hat{f}_{\boldsymbol{Y}}(\boldsymbol{y};\mathbf{H}) - f(\boldsymbol{y}) = \tfrac{1}{2}m_2(K)\operatorname{tr}(\mathbf{H}\mathsf{D}^2 f(\boldsymbol{y}))\{1+o(1)\}.$$

For the variance, we have $\operatorname{Var} \hat{f}_{\boldsymbol{Y}}(\boldsymbol{y};\mathbf{H}) = n^{-1}\,\mathbb{E}[K_{\mathbf{H}}(\boldsymbol{y}-\boldsymbol{Y})^2] - n^{-1}[\mathbb{E}\, K_{\mathbf{H}}(\boldsymbol{y}-\boldsymbol{Y})]^2$. The second term is given by the above, so we are required to evaluate

$$\begin{aligned}\mathbb{E}\, \hat{f}_{\boldsymbol{Y}}(\boldsymbol{y};\mathbf{H})^2 = \mathbb{E}[K_{\mathbf{H}}(\boldsymbol{y}-\boldsymbol{Y})^2] &= \int_{\mathbb{R}^d} K_{\mathbf{H}}(\boldsymbol{y}-\boldsymbol{w})^2 f_{\boldsymbol{Y}}(\boldsymbol{w})\, d\boldsymbol{w}\\ &= \int_{\mathbb{R}^d} |\mathbf{H}|^{-1} K(\mathbf{H}^{-1/2}(\boldsymbol{y}-\boldsymbol{w}))^2 f_{\boldsymbol{Y}}(\boldsymbol{w})\, d\boldsymbol{w}\\ &= \int_{\mathbb{R}^d} |\mathbf{H}|^{-1/2} K(\boldsymbol{w})^2 f_{\boldsymbol{Y}}(\boldsymbol{y}-\mathbf{H}^{-1/2}\boldsymbol{w}) d\boldsymbol{w}\\ &= |\mathbf{H}|^{-1/2} f_{\boldsymbol{Y}}(\boldsymbol{y}) \int_{\mathbb{R}^d} K(\boldsymbol{w})^2 d\boldsymbol{w}\{1+o(1)\}\\ &= |\mathbf{H}|^{-1/2} f_{\boldsymbol{Y}}(\boldsymbol{y}) R(K)\{1+o(1)\}.\end{aligned}$$

Thus the variance term is

$$\operatorname{Var} \hat{f}(\boldsymbol{y};\mathbf{H}) = n^{-1}\{|\mathbf{H}|^{-1/2} f_{\boldsymbol{Y}}(\boldsymbol{y})R(K) - [f_{\boldsymbol{Y}}(\boldsymbol{y}) + \tfrac{1}{2}m_2(K)\operatorname{tr}(\mathbf{H}\mathsf{D}^2 f_{\boldsymbol{Y}}(\boldsymbol{y}))]^2\}\{1+o(1)\}.$$

Since $\mathbf{H}\to 0$ then $|\mathbf{H}|^{-1/2}$ dominates both the constant term $f_{\boldsymbol{Y}}(\boldsymbol{y})$ and the $\operatorname{tr}(\mathbf{H})$ term so we can write

$$\operatorname{Var} \hat{f}_{\boldsymbol{Y}}(\boldsymbol{y};\mathbf{H}) = n^{-1}|\mathbf{H}|^{-1/2} f_{\boldsymbol{Y}}(\boldsymbol{y}) R(K)\{1+o(1)\}.$$

The integrated square bias (ISB) is then

$$\begin{aligned}\mathrm{ISB}\,\hat{f}_{\boldsymbol{Y}}(\cdot;\mathbf{H}) &= \int_{\mathbb{R}^d} \mathrm{Bias}^2\,\hat{f}_{\boldsymbol{Y}}(\boldsymbol{y};\mathbf{H})\mathrm{d}\boldsymbol{y} = \int_{\mathbb{R}^d} \tfrac{1}{4}m_2^2(K)\,\mathrm{tr}^2(\mathbf{H}\mathsf{D}^2 f(\boldsymbol{y}))\mathrm{d}\boldsymbol{y}\{1+o(1)\}\\ &= \tfrac{1}{4}m_2^2(K)\int_{\mathbb{R}^d}\mathrm{tr}^2(\mathbf{H}\mathsf{D}^2 f(\boldsymbol{y}))\mathrm{d}\boldsymbol{y}\{1+o(1)\}\\ &= \tfrac{1}{4}m_2^2(K)(\mathrm{vec}^T\,\mathbf{H}\otimes\mathrm{vec}^T\,\mathbf{H})\boldsymbol{\psi}_{\boldsymbol{Y},4}\{1+o(1)\},\end{aligned}$$

and similarly the integrated variance (IV) is

$$\begin{aligned}\mathrm{IV}\,\hat{f}_{\boldsymbol{Y}}(\cdot;\mathbf{H}) &= \int_{\mathbb{R}^d} n^{-1}|\mathbf{H}|^{-1/2} f_{\boldsymbol{Y}}(\boldsymbol{y})R(K)\mathrm{d}\boldsymbol{y}\{1+o(1)\}\\ &= n^{-1}|\mathbf{H}|^{-1/2}R(K)\int_{\mathbb{R}^d} f_{\boldsymbol{Y}}(\boldsymbol{y})\mathrm{d}\boldsymbol{y}\{1+o(1)\}\\ &= n^{-1}|\mathbf{H}|^{-1/2}R(K)\{1+o(1)\},\end{aligned}$$

using the integrability assumptions in conditions (A1) and (A2). Hence we obtain

$$\begin{aligned}\mathrm{MISE}\,\hat{f}_{\boldsymbol{Y}}(\cdot;\mathbf{H}) &= \mathrm{ISB}\,\hat{f}_{\boldsymbol{Y}}(\cdot;\mathbf{H}) + \mathrm{IV}\,\hat{f}_{\boldsymbol{Y}}(\cdot;\mathbf{H})\\ &= \big[\tfrac{1}{4}m_2^2(K)(\mathrm{vec}^T\,\mathbf{H}\otimes\mathrm{vec}^T\,\mathbf{H})\boldsymbol{\psi}_{\boldsymbol{Y},4} + n^{-1}|\mathbf{H}|^{-1/2}R(K)\big]\{1+o(1)\}.\end{aligned}$$

The optimiser of this MISE requires that the order of the squared bias and the variance are the same, and so $\mathbf{H}^* = O(n^{-2/(d+4)})$. Hence

$$\begin{aligned}\inf_{\mathbf{H}\in\mathcal{F}}\mathrm{MISE}\,\hat{f}_{\boldsymbol{Y}}(\cdot;\mathbf{H}) &= \mathrm{MISE}\,\hat{f}_{\boldsymbol{Y}}(\cdot;\mathbf{H}^*)\\ &= \big[\tfrac{1}{4}m_2^2(K)(\mathrm{vec}^T\,\mathbf{H}^*\otimes\mathrm{vec}^T\,\mathbf{H}^*)\boldsymbol{\psi}_{\boldsymbol{Y},4} + n^{-1}|\mathbf{H}^*|^{-1/2}R(K)\big]\{1+o(1)\}\\ &= O(n^{-4/(d+4)}).\end{aligned}$$

□

Proof of Theorem 1

As $\boldsymbol{X}$ is supported on $(\boldsymbol{u},\infty)$, we require the following additional assumption:

(A4) The transformation $\boldsymbol{t} = (t_1,\ldots,t_d)^\top$ is a d-vector function where each component function t_j is monotone, maps (u_j,∞) to $\mathbb{R}$ and is continuously differentiable, for $j = 1,\ldots,d$, and the Jacobian of $\boldsymbol{t}$ is square integrable.

Proof of Theorem 1. Firstly using the definition of $\hat{f}_{\boldsymbol{X}^{[\boldsymbol{u}]}}$ given in Equation (1), its expected value is $\mathbb{E}\,\hat{f}_{\boldsymbol{X}^{[\boldsymbol{u}]}}(\boldsymbol{x};\mathbf{H}) = |\mathbf{J}_{\boldsymbol{t}}(\boldsymbol{x})|\,\mathbb{E}\,\hat{f}_{\boldsymbol{Y}}(\boldsymbol{t}(\boldsymbol{x});\mathbf{H})$ and its associated bias is

$$\mathrm{Bias}\,\hat{f}_{\boldsymbol{X}^{[\boldsymbol{u}]}}(\boldsymbol{x};\mathbf{H}) = \mathbb{E}\,\hat{f}_{\boldsymbol{X}^{[\boldsymbol{u}]}}(\boldsymbol{x};\mathbf{H}) - f_{\boldsymbol{X}^{[\boldsymbol{u}]}}(\boldsymbol{x}) = |\mathbf{J}_{\boldsymbol{t}}(\boldsymbol{x})|\,\mathrm{Bias}\,\hat{f}_{\boldsymbol{Y}}(\boldsymbol{t}(\boldsymbol{x});\mathbf{H}).$$

Similarly we have $\mathbb{E}\,\hat{f}_{\boldsymbol{X}^{[\boldsymbol{u}]}}(\boldsymbol{x};\mathbf{H})^2 = |\mathbf{J}_{\boldsymbol{t}}(\boldsymbol{x})|^2\,\mathbb{E}\,\hat{f}_{\boldsymbol{Y}}(\boldsymbol{t}(\boldsymbol{x});\mathbf{H})^2$, leading to

$$\mathrm{Var}\,\hat{f}_{\boldsymbol{X}^{[\boldsymbol{u}]}}(\boldsymbol{x};\mathbf{H}) = \mathbb{E}\,\hat{f}_{\boldsymbol{X}^{[\boldsymbol{u}]}}(\boldsymbol{x};\mathbf{H})^2 - \big[\mathbb{E}\,\hat{f}_{\boldsymbol{X}^{[\boldsymbol{u}]}}(\boldsymbol{x};\mathbf{H})\big]^2 = |\mathbf{J}_{\boldsymbol{t}}(\boldsymbol{x})|^2\,\mathrm{Var}\,\hat{f}_{\boldsymbol{Y}}(\boldsymbol{t}(\boldsymbol{x});\mathbf{H}).$$

The integrated square bias (ISB) is

$$\begin{aligned}\mathrm{ISB}\,\hat{f}_{\boldsymbol{X}^{[\boldsymbol{u}]}}(\boldsymbol{x};\mathbf{H}) &= \int_{(\boldsymbol{u},\infty)} |\mathbf{J}_{\boldsymbol{t}}(\boldsymbol{x})|^2\,\mathrm{Bias}^2\,\hat{f}_{\boldsymbol{Y}}(\boldsymbol{t}(\boldsymbol{x});\mathbf{H})\mathrm{d}\boldsymbol{x}\\ &= \int_{\mathbb{R}^d} |\mathbf{J}_{\boldsymbol{t}}(\boldsymbol{t}^{-1}(\boldsymbol{y}))|^2\,\mathrm{Bias}^2\,\hat{f}_{\boldsymbol{Y}}(\boldsymbol{t}(\boldsymbol{t}^{-1}(\boldsymbol{y}));\mathbf{H})\mathrm{d}\boldsymbol{y}\end{aligned}$$

given the change of variable $\boldsymbol{y} = \boldsymbol{t}(\boldsymbol{x})$. Similarly the integrated variance (IV) is

$$\mathrm{IV}\, \hat{f}_{\boldsymbol{X}^{[\boldsymbol{u}]}}(\boldsymbol{x}; \mathbf{H}) = \int_{(\boldsymbol{u},\infty)} |\mathbf{J}_{\boldsymbol{t}}(\boldsymbol{x})|^2 \operatorname{Var} \hat{f}_{\boldsymbol{Y}}(\boldsymbol{t}(\boldsymbol{x}); \mathbf{H}) \mathrm{d}\boldsymbol{x} = \int_{\mathbb{R}^d} |\mathbf{J}_{\boldsymbol{t}}(\boldsymbol{t}^{-1}(\boldsymbol{y}))|^2 \operatorname{Var} \hat{f}_{\boldsymbol{Y}}(\boldsymbol{y}; \mathbf{H}) \mathrm{d}\boldsymbol{x}.$$

As a consequence, the MISE of $\hat{f}_{\boldsymbol{X}^{[\boldsymbol{u}]}}$ can be written as

$$\operatorname{MISE} \hat{f}_{\boldsymbol{X}^{[\boldsymbol{u}]}}(\cdot; \mathbf{H}) = \int_{\mathbb{R}^d} |\mathbf{J}_{\boldsymbol{t}}(\boldsymbol{t}^{-1}(\boldsymbol{y}))|^2 \operatorname{MSE} \hat{f}_{\boldsymbol{Y}}(\cdot; \mathbf{H}) \mathrm{d}\boldsymbol{y}.$$

Since the $|\mathbf{J}_{\boldsymbol{t}}(\boldsymbol{t}^{-1}(\boldsymbol{y}))|$ term does not involve the sample size n or the bandwidth matrix $\mathbf{H}$, it has no influence on the rate of convergence, thus the MISE of $\hat{f}_{\boldsymbol{X}^{[\boldsymbol{u}]}}$ and $\hat{f}_{\boldsymbol{Y}}$ are asymptotically equivalent. As a result

$$\operatorname{MISE} \hat{f}_{\boldsymbol{X}^{[\boldsymbol{u}]}}(\cdot; \mathbf{H}) = \{\operatorname{MISE} \hat{f}_{\boldsymbol{Y}}(\cdot; \mathbf{H})\}\{1 + o(1)\}.$$

Since $\mathbf{H}^*$ and $\mathbf{H}^{*[\boldsymbol{u}]} = \operatorname{argmin}_{\mathbf{H} \in \mathcal{F}} \operatorname{MISE} \hat{f}_{\boldsymbol{X}^{[\boldsymbol{u}]}}(\cdot; \mathbf{H})$ are of the same asymptotic $n^{-2/(d+4)}$ order, using Lemma 1 we then have

$$\begin{aligned}
&\inf_{\mathbf{H} \in \mathcal{F}} \operatorname{MISE} \hat{f}_{\boldsymbol{X}^{[\boldsymbol{u}]}}(\cdot; \mathbf{H}) - \Big\{ \inf_{\mathbf{H} \in \mathcal{F}} \operatorname{MISE} \hat{f}_{\boldsymbol{Y}}(\cdot; \mathbf{H}) \Big\} = \operatorname{MISE} \hat{f}_{\boldsymbol{X}^{[\boldsymbol{u}]}}(\cdot; \mathbf{H}^{*[\boldsymbol{u}]}) - \operatorname{MISE} \hat{f}_{\boldsymbol{Y}}(\cdot; \mathbf{H}^*) \\
&\quad = \big\{ \tfrac{1}{4} m_2^2(K) [(\operatorname{vec}^T \mathbf{H}^{*[\boldsymbol{u}]})^{\otimes 2} \boldsymbol{\psi}_{\boldsymbol{X}^{[\boldsymbol{u}]},4} - (\operatorname{vec}^T \mathbf{H}^*)^{\otimes 2} \boldsymbol{\psi}_{\boldsymbol{Y},4}] \\
&\quad + n^{-1} (|\mathbf{H}^{*[\boldsymbol{u}]}|^{-1/2} - |\mathbf{H}^*|^{-1/2}) R(K) \big\} \{1 + o(1)\} = O(n^{-4/(d+4)}).
\end{aligned}$$

□

References

Beirlant, J., Goegebeur, Y., Teugels, J., and Segers, J. (2004), *Statistics of Extremes: Theory and Applications*, Chichester: John Wiley & Sons.

Beranger, B. and Padoan, S. A. (2015), "Extreme Dependence Models," *ArXiv e-prints*.

Bowman, A. W., Hall, P., and Titterington, D. M. (1984), "Cross-validation in nonparametric estimation of probabilities and probability densities," *Biometrika*, 71, 341–351.

Chacón, J. E. and Duong, T. (2010), "Multivariate plug-in bandwidth selection with unconstrained bandwidth matrices," *Test*, 19, 375–398.

Charpentier, A. and Flachaire, E. (2014), "Log-Transform Kernel Density Estimation of Income Distribution," DOI: http://dx.doi.org/10.2139/ssrn.2514882.

Chen, S. X. (1999), "Beta kernel estimators for density functions," *Comput. Statist. Data Anal.*, 31, 131–145.

Cheng, R. C. H. and Amin, N. A. K. (1983), "Estimating parameters in continuous univariate distributions with a shifted origin," *J. Roy. Statist. Soc. Ser. B*, 45, 394–403.

Christopeit, N. (1994), "Estimating parameters of an extreme value distribution by the method of moments," *J. Statist. Plann. Inference*, 41, 173–186.

Coles, S. (2001), *An Introduction to Statistical Modeling of Extreme Values*, London: Springer-Verlag.

Coles, S. G. and Tawn, J. A. (1994), "Statistical methods for multivariate extremes: An application to structural design," *Journal of the Royal Statistical Society Series C*, 43, 1–48.

Cowan, T., Purich, A., Perkins, S., Pezza, A., Boschat, G., and Sadler, K. (2014), "More Frequent, Longer, and Hotter Heat Waves for Australia in the Twenty-First Century," *Journal of Climate*, 27, 5851–5871.

de Carvalho, M., Oumow, B., Segers, J., and Warchoł, M. (2013), "A Euclidean likelihood estimator for bivariate tail dependence," *Comm. Statist. Theory Methods*, 42, 1176–1192.

de Haan, L. and Ferreira, A. (2006), *Extreme Value Theory: An Introduction*, New York: Springer.

Devroye, L. and Györfi, L. (1985), *Nonparametric Density Estimation: the L_1 View*, New York: John Wiley and Sons.

Drees, H. and Huang, X. (1998), "Best attainable rates of convergence for estimators of the stable tail dependence function," *J. Multivariate Anal.*, 64, 25–47.

Duong, T. and Hazelton, M. L. (2003), "Plug-in bandwidth matrices for bivariate kernel density estimation," *J. Nonparametr. Stat.*, 15, 17–30.

— (2005), "Cross-validation bandwidth matrices for multivariate kernel density estimation," *Scand. J. Statist.*, 32, 485–506.

Einmahl, J. H. J., de Haan, L., and Piterbarg, V. I. (2001), "Nonparametric estimation of the spectral measure of an extreme value distribution," *Ann. Statist.*, 29, 1401–1423.

Einmahl, J. H. J., Krajina, A., and Segers, J. (2008), "A method of moments estimator of tail dependence," *Bernoulli*, 14, 1003–1026.

— (2012), "An M-estimator for tail dependence in arbitrary dimensions," *Ann. Statist.*, 40, 1764–1793.

Einmahl, J. H. J. and Segers, J. (2009), "Maximum empirical likelihood estimation of the spectral measure of an extreme-value distribution," *Ann. Statist.*, 37, 2953–2989.

Falk, M., Hüsler, J., and Reiss, R.-D. (2011), *Laws of Small Numbers: Extremes and Rare Events*, Basel: Birkhäuser/Springer Basel AG, extended ed.

Fischer, E. M., Beyerle, U., and Knutti, R. (2013), "Robust spatially aggregated projections of climate extremes," *Nature Clim. Change*, 3, 1033–1038.

Flato, G., Marotzke, J., Abiodun, B., Braconnot, P., Chou, S., Collins, W., Cox, P., Driouech, F., Emori, S., Eyring, V., Forest, C., Gleckler, P., Guilyardi, E., Jakob, C., Kattsov, V., Reason, C., and Rummukainen, M. (2013), *Evaluation of Climate Models*, Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press, book section 9, pp. 741–866.

Gasser, T. and Müller, H.-G. (1979), "Kernel estimation of regression functions," in *Smoothing Techniques for Curve Estimation*, eds. Gasser, T. and Rosenblatt, M., Berlin: Springer, pp. 23–68.

Geenens, G. (2014), "Probit Transformation for Kernel Density Estimation on the Unit Interval," *J. Amer. Statist. Assoc.*, 109, 346–358.

Hall, P., Marron, J. S., and Park, B. U. (1992), "Smoothed cross-validation," *Probab. Theory Related Fields*, 92, 1–20.

Hall, P. and Tajvidi, N. (2000), "Distribution and dependence-function estimation for bivariate extreme-value distributions," *Bernoulli*, 6, 835–844.

Hill, B. M. (1975), "A simple general approach to inference about the tail of a distribution," *Ann. Statist.*, 3, 1163–1174.

Holmquist, B. (1996), "The d-variate vector Hermite polynomial of order k," *Linear Algebra and Its Applications*, 237/238, 155–190.

Hosking, J. R. M. (1985), "Algorithm AS 215: Maximum-Likelihood Estimation of the Parameters of the Generalized Extreme-Value Distribution," *J. Roy. Stat. Soc. Ser. C*, 34, pp. 301–310.

Hosking, J. R. M., Wallis, J. R., and Wood, E. F. (1985), "Estimation of the generalized extreme-value distribution by the method of probability-weighted moments," *Technometrics*, 27, 251–261.

Huang, X. (1992), "Statistics of Bivariate Extreme Values," Ph.D. thesis, Erasmus University.

Hüsler, J. and Reiss, R.-D. (1989), "Maxima of normal random vectors: between independence and complete dependence," *Statist. Probab. Lett.*, 7, 283–286.

Joe, H. (1990), "Families of min-stable multivariate exponential and multivariate extreme value distributions," *Statist. Probab. Lett.*, 9, 75–81.

Jones, D. A., Wang, W., and Fawcett, R. (2009), "High-quality spatial climate data-sets for Australia," *Australian Meteorological and Oceanographic Journal*, 58, 233–1026.

Kotz, S. and Nadarajah, S. (2000), *Extreme Value Distributions: Theory and Applications*, London: Imperial College Press.

Lye, L., Hapuarachchi, K., and Ryan, S. (1993), "Bayes estimation of the extreme-value reliability function," *IEEE Trans. Reliab.*, 42, 641–644.

Macleod, A. J. (1989), "Remark AS R76: A Remark on Algorithm AS 215: Maximum-Likelihood Estimation of the Parameters of the Generalized Extreme-Value Distribution," *J. Roy. Statist. Soc. Ser. C*, 38, 198–199.

Marcon, G., Padoan, S. A., Naveau, P., and Muliere, P. (2014), "Multivariate nonparametric estimation of the pickands dependence function using Bernstein polynomials," *arXiv:1405.5228*.

Maritz, J. S. and Munro, A. H. (1967), "On the use of the generalised extreme-value distribution in estimating extreme percentiles," *Biometrics*, 23, 79–103.

Markovich, N. (2007), *Nonparametric Analysis of Univariate Heavy-tailed Data: Research and Practice*, Chichester: John Wiley & Sons.

Marron, J. and Ruppert, D. (1994), "Transformation to Reduce Boundary Bias in Kernel Density-Estimation," *Journal of the Royal Statistical Society Series B*, 56, 653–671.

Marshall, J. C. and Hazelton, M. L. (2010), "Boundary kernels for adaptive density estimators on regions with irregular boundaries," *J. Multivariate Anal.*, 101, 949–963.

Perkins, S. E., Moise, A., Whetton, P., and Katzfey, J. (2014), "Regional changes of climate extremes over Australia - a comparison of regional dynamical downscaling and global climate model simulations," *Int. J. Climatol.*, 34, 3456–3478.

Perkins, S. E., Pitman, A. J., Holbrook, N. J., and McAneney, J. (2007), "Evaluation of the AR4 Climate Models' Simulated Daily Maximum Temperature, Minimum Temperature, and Precipitation over Australia Using Probability Density Functions," *J. Climate*, 20, 4356–4376.

Perkins, S. E., Pitman, A. J., and Sisson, S. A. (2013), "Systematic differences in future 20 year temperature extremes in AR4 model projections over Australia as a function of model skill," *Int. J. Climatol.*, 33, 1153–1167.

Pickands, III, J. (1975), "Statistical inference using extreme order statistics," *Ann. Statist.*, 3, 119–131.

Prescott, P. and Walden, A. T. (1980), "Maximum likelihood estimation of the parameters of the generalized extreme-value distribution," *Biometrika*, 67, 723–724.

Rudemo, M. (1982), "Empirical choice of histograms and kernel density estimators," *Scand. J. Statist.*, 9, 65–78.

Ruppert, D. and Cline, D. B. H. (1994), "Transformation kernel density estimation – bias reduction by empirical transformations," *Ann. Statist.*, 22, 185–210.

Ruppert, D. and Wand, M. P. (1992), "Correcting for kurtosis in density estimation," *Austral. J. Statist.*, 34, 19–29.

Scott, D. W. (2015), *Multivariate Density Estimation: Theory, Practice, and Visualization*, Hoboken, NJ: John Wiley & Sons, 2nd ed.

Seneviratne, S., Nicholls, N., Easterling, D., Goodess, C., Kanae, S., Kossin, J., Luo, Y., Marengo, J., McInnes, K., Rahimi, M., Reichstein, M., Sorteberg, A., Vera, C., and Zhang, X. (2012), *Changes in Climate Extremes and their Impacts on the Natural Physical Environment*, Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press, book section 3, pp. 190–230.

Sheather, S. J. and Jones, M. C. (1991), "A reliable data-based bandwidth selection method for kernel density estimation," *J. Roy. Statist. Soc. Ser. B*, 53, 683–690.

Sillmann, J., Kharin, V. V., Zhang, X., Zwiers, F. W., and Bronaugh, D. (2013a), "Climate extremes indices in the CMIP5 multimodel ensemble: Part 1. Model evaluation in the present climate," *J. Geophys. Res.-Atmos.*, 118, 1716–1733.

Sillmann, J., Kharin, V. V., Zwiers, F. W., Zhang, X., and Bronaugh, D. (2013b), "Climate extremes indices in the CMIP5 multimodel ensemble: Part 2. Future climate projections," *J. Geophys. Res.-Atmos.*, 118, 2473–2493.

Silverman, B. W. (1986), *Density Estimation for Statistics and Data Analysis*, London: Chapman and Hall.

Simonoff, J. S. (1996), *Smoothing Methods in Statistics*, New York: Springer-Verlag.

Smith, R. L. (1985), "Maximum likelihood estimation in a class of nonregular cases," *Biometrika*, 72, 67–90.

Smith, R. L., Tawn, J. A., and Yuen, H. K. (1990), "Statistics of Multivariate Extremes," *Int. Statist. Rev.*, 58, 47–58.

Taylor, K. E., Stouffer, R. J., and Meehl, G. A. (2012), "An Overview of CMIP5 and the Experiment Design," *Bulletin of the American Meteorological Society*, 93, 485–498.

Vautard, R., Gobiet, A., Jacob, D., Belda, M., Colette, A., Déqué, M., Fernández, J., García-Díez, M., Goergen, K., Güttler, I., Halenka, T., Karacostas, T., Katragkou, E., Keuler, K., Kotlarski, S., Mayer, S., van Meijgaard, E., Nikulin, G., Patarcić, M., Scinocca, J., Sobolowski, S., Suklitsch, M., Teichmann, C., Warrach-Sagi, K., Wulfmeyer, V., and Yiou, P. (2013), "The simulation of European heat waves from an ensemble of regional climate models within the EURO-CORDEX project," *Climate Dynamics*, 41, 2555–2575.

Wand, M., Marron, J., and Ruppert, D. (1991), "Transformations in Density-Estimation," *Journal of the American Statistical Association*, 86, 343–353.

Wand, M. P. (1992), "Error analysis for general multivariate kernel estimators," *J. Nonparamtr. Statist.*, 2, 1–15.

Wand, M. P. and Jones, M. C. (1994), "Multivariate plug-in bandwidth selection," *Comput. Statist.*, 9, 97–116.

— (1995), *Kernel Smoothing*, London: Chapman and Hall.

Supplementary Material for 'Exploratory data analysis for extreme values using non-parametric kernel methods'

B. Béranger, T. Duong, S. E. Perkins-Kirkpatrick and S. A. Sisson

Abstract

This document contains similar results to the univariate analyses in Section 3.1 (i.e. Figures 1 and 2, and Table 1) except that the tail sample size is $n = 25$ and $n = 50$. The main manuscript uses $n = 100$.

A Simulated data - Univariate: tail sample sizes $n = 25$ and 50

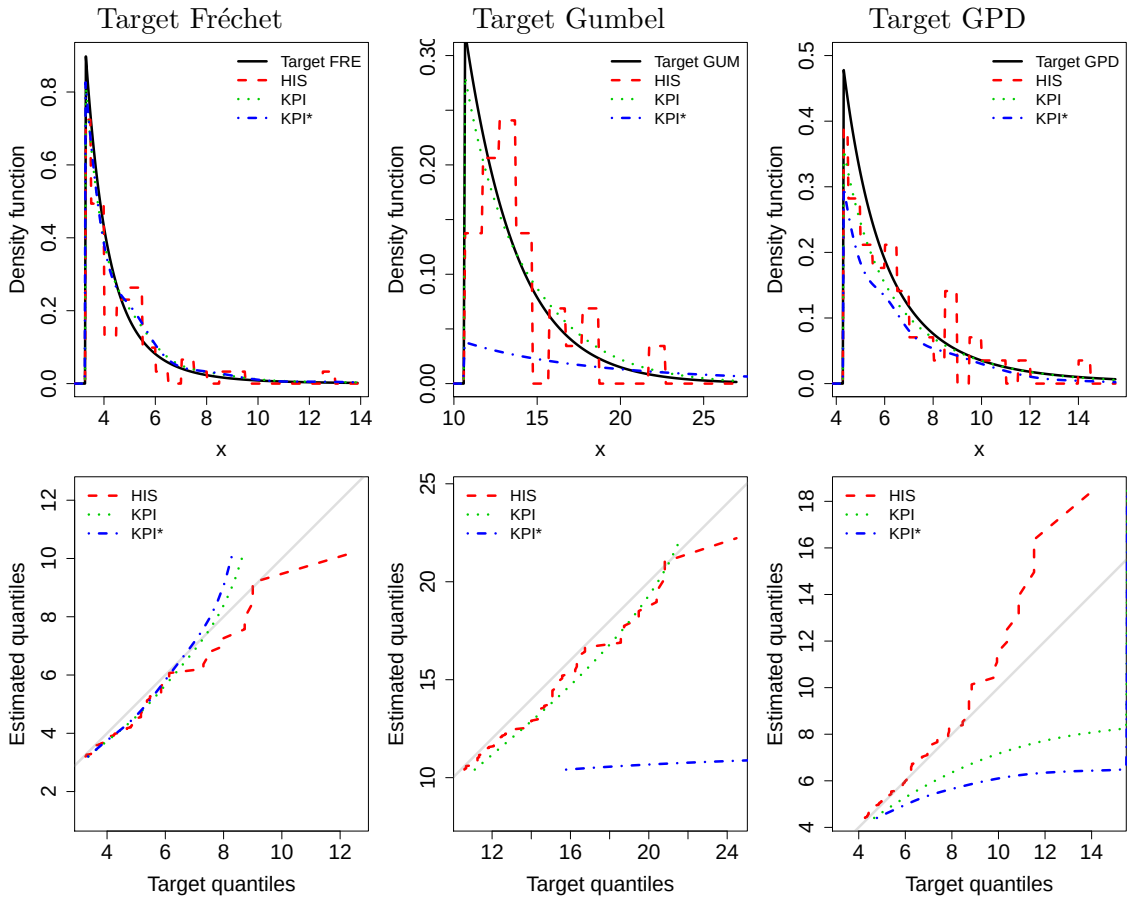

Figure 7: Non-parametric estimators of the univariate tail density (top) and of the tail quantiles (bottom) when the target density is Fréchet (left), Gumbel (centre) and generalised Pareto (right). Sample size is $n = 50$. Fréchet ($\mu = 1$, $\sigma = 0.5$, $\xi = 0.25$), Gumbel ($\mu = 1.5$, $\sigma = 3$) and Pareto ($\mu = 0$, $\sigma = 1$, $\xi = 0.25$) target densities are represented by a solid black line. The histogram estimator $\tilde{f}_{X^{[u]}}$ with normal scale binwidth (HIS) is represented by a dashed red line, the transformed kernel plug-in estimator $\hat{f}_{X^{[u]}}$ (KPI) by a dotted green line and the standard kernel estimator $\hat{f}^*_{X^{[u]}}$ (KPI*) by a dot-dash blue line.

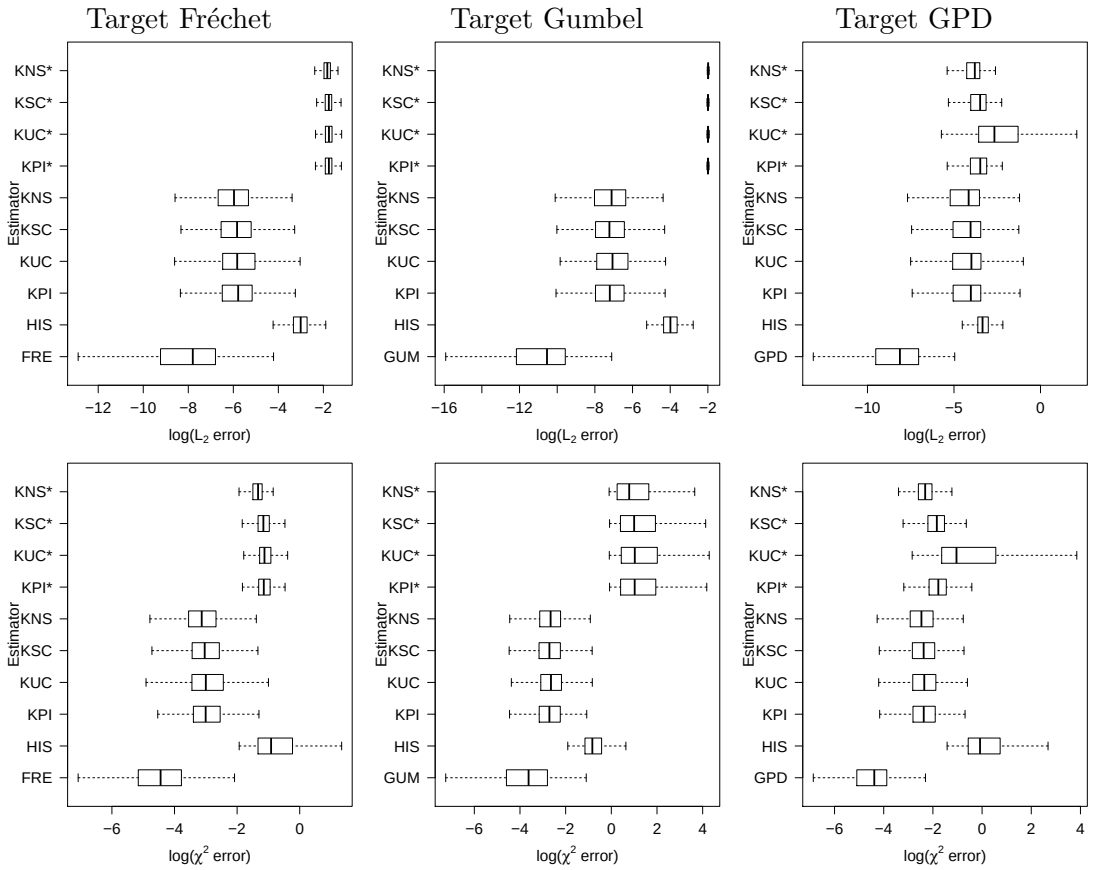

Figure 8: Box-plots of the $\log L_2$ (top panels) and $\log \chi^2$ (bottom) errors for the parametric Fréchet (FRE), Gumbel (GUM), generalised Pareto (GPD), and histogram (HIS) tail density estimators. Transformed kernel density estimators $\hat{f}_{X^{[u]}}$ use the plug-in (KPI), unbiased cross validation (KUC), smoothed cross validation (KSC) and normal scale kernel (KNS) optimal bandwidth selectors. Standard kernel density estimators $\hat{f}^*_{X^{[u]}}$ are indicated by an asterisk (*). True target densities are (left panels) Fréchet, (centre) Gumbel and (right) GPD. Box plots are based on 400 replicates of $m = 1,000$ observations with a tail sample size of $n = 50$.

		$\tilde{T}$ selected model			$\hat{T}$ selected model			$\hat{T}^*$ selected model		
Error	Target	FRE	GUM	GPD	FRE	GUM	GPD	FRE	GUM	GPD
	FRE	**0.15**	0.28	0.56	**0.90**	0.00	0.10	**0.88**	0.00	0.12
L_2	GUM	0.01	**0.76**	0.22	0.05	**0.91**	0.04	0.02	**0.98**	0.00
	GPD	0.03	0.64	**0.33**	0.14	0.03	**0.83**	0.02	0.24	**0.74**
	FRE	**1.00**	0.00	0.00	**1.00**	0.00	0.00	**1.00**	0.00	0.00
χ^2	GUM	0.05	**0.86**	0.09	0.06	**0.94**	0.00	0.06	**0.94**	0.00
	GPD	0.81	0.00	**0.19**	0.42	0.00	**0.58**	0.66	0.00	**0.34**

Table 5: Proportion of 400 simulated datasets with tail sample size $n = 50$ from each known target distribution (Fréchet, Gumbel and GPD) that are identified as coming from each of these distributions by having the smallest tail index value, as a function of nonparametric density estimator. Bold text indicates the proportion selecting the correct model. Nonparametric density estimators are the histogram ($\tilde{T}$), the transformed kernel ($\hat{T}$) and the standard kernel ($\hat{T}^*$), with tail indices calculated according to the L_2 (top row) and χ^2 loss (bottom row).

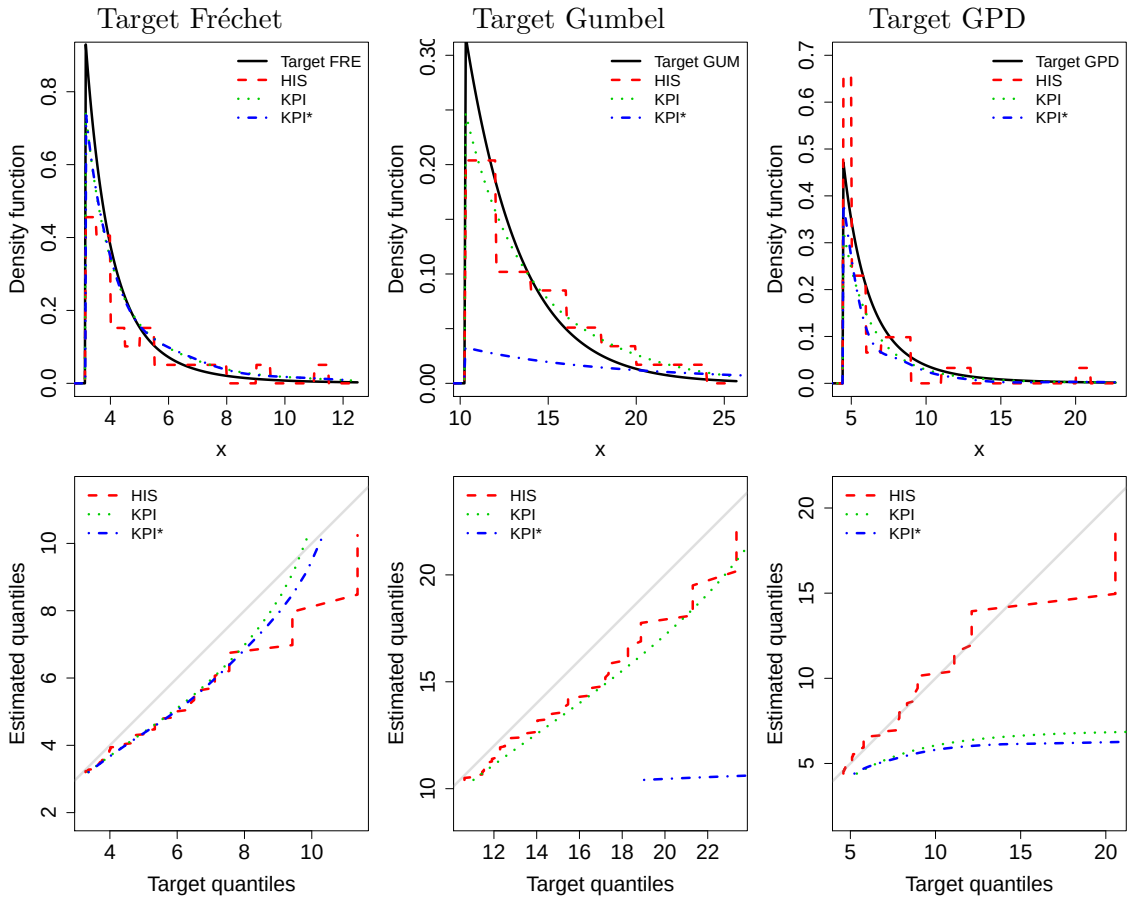

Figure 9: Non-parametric estimators of the univariate tail density (top) and of the tail quantiles (bottom) when the target density is Fréchet (left), Gumbel (centre) and generalised Pareto (right). Sample size is $n = 25$. Fréchet ($\mu = 1$, $\sigma = 0.5$, $\xi = 0.25$), Gumbel ($\mu = 1.5$, $\sigma = 3$) and Pareto ($\mu = 0$, $\sigma = 1$, $\xi = 0.25$) target densities are represented by a solid black line. The histogram estimator $\tilde{f}_{X^{[u]}}$ with normal scale binwidth (HIS) is represented by a dashed red line, the transformed kernel plug-in estimator $\hat{f}_{X^{[u]}}$ (KPI) by a dotted green line and the standard kernel estimator $\hat{f}^*_{X^{[u]}}$ (KPI*) by a dot-dash blue line.

		$\tilde{T}$ selected model			$\hat{T}$ selected model			$\hat{T}^*$ selected model		
Error	Target	FRE	GUM	GPD	FRE	GUM	GPD	FRE	GUM	GPD
L_2	FRE	**0.14**	0.49	0.37	**0.90**	0.00	0.10	**0.85**	0.00	0.15
	GUM	0.02	**0.74**	0.24	0.06	**0.87**	0.07	0.01	**0.96**	0.02
	GPD	0.06	0.73	**0.20**	0.26	0.02	**0.72**	0.13	0.20	**0.68**
χ^2	FRE	**1.00**	0.00	0.00	**1.00**	0.00	0.00	**1.00**	0.00	0.00
	GUM	0.06	**0.86**	0.08	0.06	**0.94**	0.00	0.06	**0.94**	0.00
	GPD	0.77	0.00	**0.23**	0.50	0.00	**0.50**	0.62	0.00	**0.38**

Table 6: Proportion of 400 simulated datasets with tail sample size $n = 25$ from each known target distribution (Fréchet, Gumbel and GPD) that are identified as coming from each of these distributions by having the smallest tail index value, as a function of nonparametric density estimator. Bold text indicates the proportion selecting the correct model. Nonparametric density estimators are the histogram ($\tilde{T}$), the transformed kernel ($\hat{T}$) and the standard kernel ($\hat{T}^*$), with tail indices calculated according to the L_2 (top row) and χ^2 loss (bottom row).

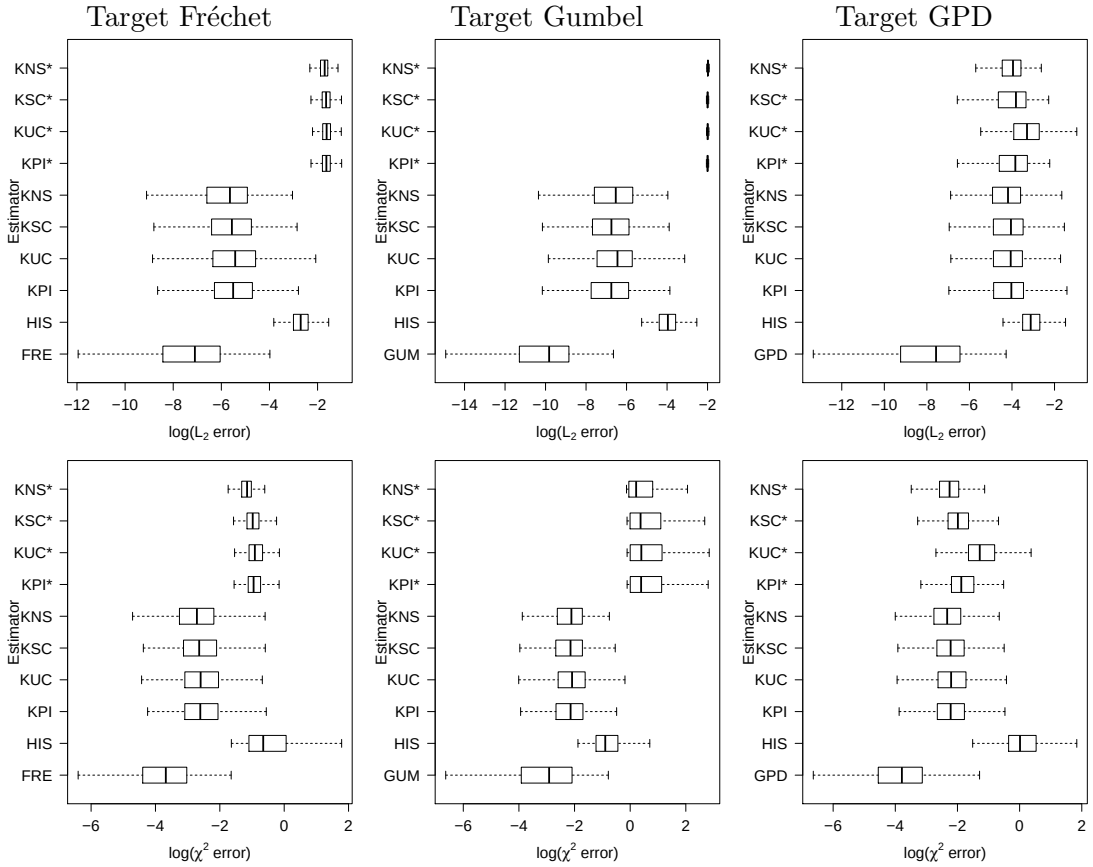

Figure 10: Box-plots of the $\log L_2$ (top panels) and $\log \chi^2$ (bottom) errors for the parametric Fréchet (FRE), Gumbel (GUM), generalised Pareto (GPD), and histogram (HIS) tail density estimators. Transformed kernel density estimators $\hat{f}_{X^{[u]}}$ use the plug-in (KPI), unbiased cross validation (KUC), smoothed cross validation (KSC) and normal scale kernel (KNS) optimal bandwidth selectors. Standard kernel density estimators $\hat{f}^*_{X^{[u]}}$ are indicated by an asterisk (*). True target densities are (left panels) Fréchet, (centre) Gumbel and (right) GPD. Box plots are based on 400 replicates of $m = 500$ observations with a tail sample size of $n = 25$.